

DIGITALES ARCHIV

ZBW – Leibniz-Informationszentrum Wirtschaft
ZBW – Leibniz Information Centre for Economics

Elsner, Wolfram (Ed.)

Book

Game theory and institutional economics

Reference: (2014). Game theory and institutional economics. 1. ed.. Basel [u.a.] : MDPI.

This Version is available at:
<http://hdl.handle.net/11159/65>

Kontakt/Contact

ZBW – Leibniz-Informationszentrum Wirtschaft/Leibniz Information Centre for Economics
Düsternbrooker Weg 120
24105 Kiel (Germany)
E-Mail: [rights\[at\]zbw.eu](mailto:rights[at]zbw.eu)
<https://www.zbw.eu/econis-archiv/>

Standard-Nutzungsbedingungen:

Dieses Dokument darf zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden. Sie dürfen dieses Dokument nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen. Sofern für das Dokument eine Open-Content-Lizenz verwendet wurde, so gelten abweichend von diesen Nutzungsbedingungen die in der Lizenz gewährten Nutzungsrechte.

<https://zbw.eu/econis-archiv/termsfuse>

Terms of use:

This document may be saved and copied for your personal and scholarly purposes. You are not to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute or otherwise use the document in public. If the document is made available under a Creative Commons Licence you may exercise further usage rights as specified in the licence.



Game Theory and Institutional Economics

Edited by

Wolfram Elsner, Torsten Heinrich, Henning Schwardt
and Claudius Gräbner

Printed Edition of the Special Issue Published in *Games*



www.mdpi.com/journal/games

Wolfram Elsner, Torsten Heinrich,
Henning Schwardt and Claudius Gräbner (Eds.)

Game Theory and Institutional Economics



This book is a reprint of the special issue that appeared in the online open access journal *Games* (ISSN 2073-4336) in 2014 (available at: http://www.mdpi.com/journal/games/special_issues/game_theory_instit_econom).

Guest Editors

Wolfram Elsner, Torsten Heinrich, Henning Schwardt, Claudius Gräbner
Institute for Institutional and Innovation Economics (iino),
Department of Business Studies and Economics,
University of Bremen, Bremen, Germany

Editorial Office

MDPI AG
Klybeckstrasse 64
Basel, Switzerland

Publisher

Shu-Kun Lin

Managing Editor

Alistair Freeland

1. Edition 2014

MDPI • Basel • Beijing

ISBN 978-3-03842-014-9

© 2014 by the authors; licensee MDPI, Basel, Switzerland. All articles in this volume are Open Access distributed under the Creative Commons Attribution 3.0 license (<http://creativecommons.org/licenses/by/3.0/>), which allows users to download, copy and build upon published articles even for commercial purposes, as long as the author and publisher are properly credited, which ensures maximum dissemination and a wider impact of our publications. However, the dissemination and distribution of copies of this book as a whole is restricted to MDPI, Basel, Switzerland.

Table of Contents

List of Contributors	V
Wolfram Elsner, Torsten Heinrich, Henning Schwardt and Claudius Gräbner Special Issue: Aspects of Game Theory and Institutional Economics—Editorial	
<i>Games</i> 2014 , 5(3), 188-190	1
http://www.mdpi.com/2073-4336/5/3/188	
Manuel Wäckerle, Bernhard Rengs and Wolfgang Radax An Agent-Based Model of Institutional Life-Cycles	
<i>Games</i> 2014 , 5(3), 160-187	4
http://www.mdpi.com/2073-4336/5/3/160	
Jürgen Fleiß and Stefan Palan Of Coordinators and Dictators: A Public Goods Experiment	
<i>Games</i> 2014 , 4(4), 584-607	33
http://www.mdpi.com/2073-4336/4/4/584	
Tassos Patokos Introducing Disappointment Dynamics and Comparing Behaviors in Evolutionary Games: Some Simulation Results	
<i>Games</i> 2013 , 5(1), 1-25	55
http://www.mdpi.com/2073-4336/5/1/1	
Alexander J. Field Schelling, von Neumann, and the Event that Didn't Occur	
<i>Games</i> 2013 , 5(1), 53-89.....	78
http://www.mdpi.com/2073-4336/5/1/53	

List of Contributors

Manuel Wäckerle

Institute for Ecological Economics, Department Socioeconomics, Vienna University of Economics and Business, Welthandelsplatz 1/D4/2nd Floor/D4.2.232, 1020 Wien, Austria

Bernhard Rengs

Institute of Mathematical Methods in Economics, Research Group Economics, Vienna University of Technology, Argentinierstraße 8/4/105-3, 1040 Wien, Austria

Wolfgang Radax

Institute of Mathematical Methods in Economics, Research Group Economics, Vienna University of Technology, Argentinierstraße 8/4/105-3, 1040 Wien, Austria

Jürgen Fleiß

Institute of Statistics and Operations Research, University of Graz, Universitätsstrasse 15, 8010 Graz, Austria

Stefan Palan

Department of Banking and Finance, University of Innsbruck, Universitätsstrasse 15, 6020 Innsbruck, Austria; Institute of Banking and Finance, University of Graz, Universitätsstrasse 15/F2, 8010 Graz, Austria

Tassos Patokos

University of Hertfordshire, Hertfordshire Business School, Department of Accounting, Finance and Economics, Hatfield, AL10 9AB, UK

Alexander J. Field

Department of Economics, Santa Clara University, Santa Clara, CA 95053, USA

Special Issue: Aspects of Game Theory and Institutional Economics—Editorial

Wolfram Elsner, Torsten Heinrich, Henning Schwardt and Claudius Gräbner

Reprinted from Special Issue: Aspects of Game Theory and Institutional Economics, *Games*. Cite as: Elsner, W.; Heinrich, T.; Schwardt, H.; Gräbner, C. Special Issue: Aspects of Game Theory and Institutional Economics (Editorial). *Games* **2014**, *5*, 188–190.

1. Towards a Complexity Economics

Classical economists from Adam Smith to Thomas Malthus and to Karl Marx have considered the importance of *direct interdependence* and direct interactions for the economy. This was even more the case for original institutionalist thinkers such as Thorstein Veblen, John Commons, and Clarence Ayres. In their writings, direct interdependence, interactions (or transactions) among agents, with all beneficial and with all problematic consequences, took center stage in economic analysis. Why, for instance, do people adhere to a particular new fashion or trend? Because others do, after eminent people, wealthy people, the “leisure class” (T. Veblen), have made it a symbol for status. The new fashion, however, ceases to serve as such a symbol once too many people follow it. The constant effort put into following trends and adopting fashion turns out to be a *social dilemma*, driven by Veblenian instincts, such as invidious distinction in predatory societies, conspicuous consumption and emulation.

The general issues of herd behavior and myopic individualistic decision making, both under opacity and highly bounded rationality in complex systems, and both possibly carrying the problem of *negative unintended consequences* for the economy as a whole, for society and the natural commons, have taken center stage again in the global financial crisis of 2007, which still lingers around as the “Great Recession”.

The “*representative agent*” of the economic mainstream’s theoretical core model, whose decisions may simplistically be summed up into the aggregates of the related macroeconomics, has most overtly failed to enable economists to even realize inherent tendencies towards crises in real-world complex systems, based on intricate common and collective decision structures.

Of course, both classical political economy and original institutional economics lacked the formal methods to describe intricate interdependencies and decision structures in exact and mathematical terms—even if they wanted to (while some of them have been rather critical with respect to any formalism). In particular, they could not employ modern *game theory*, *evolutionary algorithms*, or *agent-based simulation*. Today, we can. As it happens, these methods are much better suited to describe classical and evolutionary-institutionalist theories and complex economies of interactive agents than, for instance, vector fields and nonlinear optimization that were *en vogue* in economics at the time of Veblen’s writings. The ex-post aggregates that emerge in simulations show one instance in which analytical advantages can in fact be realized utilizing such methods.

Among economists, the interest in these approaches has continued to grow. Making classical and evolutionary-institutionalist theories accessible to formal methods, and in this way even shedding

new light on different aspects of traditional subjects has been at the core of a new development in economics—using game theory, agent-based modelling, simulation, and lab experiments to tackle complex dynamic, evolutionary, and institutional phenomena. *Elinor Ostrom*, well-known for her experimental as well as theoretical and field research on the commons and social dilemmas, received the Nobel memorial prize in 2009, but many other scholars (including a number of other Nobel laureates) have been active in this field as well. In fact, such *complexity economics* has apparently become a new vanishing point of the economics discipline, at least in research (albeit not in textbooks and academic mass education).

2. The Contributions to this Special Issue

The current Special Issue of “games” is an attempt to highlight some recent work in this field and to bring some papers of the field together in a single publication.

While two of the papers (by Wolfgang Radax, Bernhard Rengs, and Manuel Wäckerle, and that of Jürgen Fleiß and Stefan Palan) in this issue pursue the question of the *emergence and coevolution of institutions and hierarchy*, a third paper (by Tassos Patokos) analyzes algorithms of strategy change in evolutionary game models, and the fourth paper (by Alexander Field) takes a historical (history-of-economic-thought) point of view on the development of game theory during the cold war.

Radax et al. (“An Agent-Based Model of Institutional Life-Cycles”) offer an application of an agent-based simulation on the *formation and development of institutions*. They consider the development of institutions in a population of agents playing repeated *prisoners’ dilemma* games. The paper models institutions as voluntary associations of agents with cooperation enforced by a leader. Simulations show three sharply differing possible scenarios: universal institutionalized cooperation, universal defection with stable institutions of internal cooperation, and a scenario with non-trivial *institutional life-cycles*. In an elegant way, the authors connect the paper metaphorically and practically to Ilya Prigogine’s and Stuart Kauffman’s concepts of complexity as a *transition regime between ordered and chaotic states*, identifying these three regimes, each with one of the dominant three scenarios they found in their simulation study.

Fleiß and Palan (“Of Coordinators and Dictators: A Public Goods Experiment”) consider a similar question with a methodologically different and, in fact, complementary approach. They conducted *laboratory experiments* and find that human subjects are generally willing—at an overwhelming margin—to pay for being part of an *institution* with enforced cooperation when faced with a *social dilemma* situation. The context was the production of a public good; agents could freely choose between a setting with voluntary and one with enforced contribution. In their experiments, agents strongly favor *enforced contribution* even if the randomly selected leader of the institution can exploit her position and free-ride.

Patokos (“Introducing Disappointment Dynamics and Comparing Behaviors in Evolutionary Games: Some Simulation Results”) pursues a subject that is of central importance to many *evolutionary game theory* and *replicator models* (those with a reassignment of strategies instead of exit and entry): the mode of *strategy updating* on the part of the agents. He considers three commonly used algorithms, immediate updating in the event of sub-optimal outcomes, updating based on the outcomes of several iterations, and updating based on a threshold outcome-level. Analyzing the

algorithms against a *number of game structures* for evolutionary game theory interaction settings and again with stochastic perturbation, sharply diverging outcomes are found.

Field (“Schelling, von Neumann, and the Event that Didn’t Occur”), in a quite different kind of paper, a historical review, discusses the evolution of game theory during the cold war. He argues that the historical standoff confrontation at the edge of a nuclear war had a profound impact on, not only the scholars, but also the ideas, concepts, and methods of game theory itself. While this historical reconstruction of a critical phase of development, application, and identity finding of game theory, may easily be controversial, it offers a number of truly challenging thoughts and reflections that are as original as they are unorthodox—provided, in this case, by an established institutional(ist) game theorist.

3. An Outlook

The study of institutions is not a new subject in economics at all. But it has attracted a continuing scholarly interest for decades, and many conceptual, theoretical, and methodological breakthroughs have been accomplished in this field, particularly since the dawn of modern game theory and the use of computers for simulation and economic experiments. The seminal works by Robert Axelrod and Elinor Ostrom are merely two examples. These too were centered on the *evolution of institutions in a context of social dilemma situations*—a research project that is vigorously continued by the first two papers in this issue. Of course, these efforts must always be accompanied by an equally firm resolution to work on other questions of evolutionary and institutional economics (and beyond); further on methodology and, finally, to critically question the institutional history of both game theory formal methods and evolutionary and institutional economics (on which this issue also contains one paper each).

Bremen, Germany, May 2014

An Agent-Based Model of Institutional Life-Cycles

Manuel Wäckerle, Bernhard Rengs and Wolfgang Radax

Abstract: We use an agent-based model to investigate the interdependent dynamics between individual agency and emergent socioeconomic structure, leading to institutional change in a generic way. Our model simulates the emergence and exit of institutional units, understood as generic governed social structures. We show how endogenized trust and exogenously given leader authority influences institutional change, *i.e.*, diversity in institutional life-cycles. It turns out that these governed institutions (de)structure in cyclical patterns dependent on the overall evolution of trust in the artificial society, while at the same time, influencing this evolution by supporting social learning. Simulation results indicate three scenarios of institutional life-cycles. Institutions may, (1) build up very fast and freeze the artificial society in a stable but fearful pattern (ordered system); (2) exist only for a short time, leading to a very trusty society (highly fluctuating system); and (3) structure in cyclical patterns over time and support social learning due to cumulative causation of societal trust (complex system).

Reprinted from Special Issue: Aspects of Game Theory and Institutional Economics, *Games*. Cite as: Wäckerle, M.; Rengs, B.; Radax, W. An Agent-Based Model of Institutional Life-Cycles. *Games* **2014**, *5*, 160–187.

1. Introduction

The central research question of this paper deals with the emergent effects originating from dynamic interdependencies of individual strategies and social structures—see [1] for the latter. The theory of games originally established by von Neumann and Morgenstern is perfectly suited to approach such a problem, since it provides a formal mathematical body to model social interaction and basic communication structures. Thus far, game theory was used to model a multitude of socioeconomic problems, assigning relevance to strategy formation as a major influence on economic behavior, see [2] for a recent overview on the integration of game theory and the behavioral sciences. However, with the rise of evolutionary game theory the notion of learning became a central issue of investigation, in particular within population dynamics, see [3] for one of the first elaborations. In this context population dynamics have become central in phylogenetic terms, where evolutionary stable strategies (ESS) enhanced the fitness of a group compared to others within a population. The evolutionary turn has revealed a very important finding, namely that certain strategies of conditional cooperation may lead to an ESS, thereby outplaying the strict dominance of the defective strategy, thus, transforming non-cooperative games into cooperative games, see [4,5]. Obviously this strand of research has also influenced findings about our own origin and heritage as a human species, see [6]. Although we have gathered tremendous knowledge on our social preferences and the strategic sources for cooperation in non-kin large-scale societies, we have not properly connected these findings yet with the emergence, life and exit of institutions in economy and society, *i.e.*, institutional change. This research topic opens up a multitude of interesting research questions, first

attempts to cope with them are given in [7]. Furthermore Ostrom [8] has clarified that institutions coordinate individual strategies concerning collective action problems. They represent more than just constraints on behavior [9], but may even lead to the emergence of new forms of behavior. Institutions evolve along strategies and rules, in consequence, they evolve in diverse forms from social interaction and reconstitute economic behavior. Seemingly, it stands to reason, that institutions are meta-stabilized sets of established and culturally transmitted [10] rules forming the cornerstones of political economy and its evolution. In this perspective, the theory of games can play a decisive role in explaining the political economic causes of endogenous crisis [11] via the accumulation of historically established strategies, habits, and their potential lock-in resulting in unequal patterns of social stratification, see [12] for Bourdieu's sociological analysis of the problem at hand. A similar socioeconomic approach got established almost 100 years earlier by Thorstein Veblen [13]. Veblen has looked into the cumulative causation of habits of thought resulting in institutional change. This first socio-cybernetic approach was interpreted by him as an evolutionary contribution to economics [14], because evolving institutions depend on the variation, selection and retention of habits of thought and social norms, such as conspicuous consumption for instance. Thereby, institutions are understood as social structures, which, again, feed-back to the establishment of new habits and norms. These spiral dynamics are crucial for what has been called the old institutional economics. The old approach to institutionalism stands in contrast to the new institutional economics research program where attention is turned to the transaction costs of socioeconomic activities, compare [15] for the demarcation problem between old and new institutional economics. Obviously, it is the notion of contingent path-dependent evolution, which makes the former approach richer in scope but more difficult to model. However, today we have the analytical tools to compete with such a challenge, as the original attempt by [16] has recently shown.

In order to fully integrate the theory of games into an evolutionary approach of institutional change as the central sub-field of evolutionary political economy, [11] suggests considering a computational and algorithmic methodology of agent-based modeling (ABM) and socioeconomic simulation. Recently this attempt has received increasing attention within a certain part of evolutionary economics, concerned with institutional evolution, see [17] for a computational multi-agent approach to meso-economics and critical platform size. The ABM approach suits the problem at hand well, because it is able to mimic the complex non-equilibrium dynamics of an evolving economy. The analyzed emergent properties are revealed on a meso-economic level, between micro and macro [18], acknowledged recently by Arthur [19]. Institutions play a central role in this process as social structures of rule correspondence, however, in the history of economic thought, heuristics were always explained differently, compare [20]. The advantage of the agent-based methodology over evolutionary game theory is given by the possibility to model institutions as accumulating social structures, once certain rules are introduced about governance and regulation. From a static analytical perspective, Aoki [21] has provided the first theoretical framework to analyze institutional complementarities via one-shot games as strategic systems of shared beliefs. Still, in this realm of research, dynamic models of interdependent agency-structure relations causing institutional change are rare and need further attention.

2. Model

We propose a framework to model the emergence, life and exit of institutions (institutional life cycles) in an artificial political economy based on the interactions of individuals on a micro level. In the model we treat institutions as social accumulating structures instead of mere sets of agents with common properties, where the frequency and coordination of strategies and behavioral motivations plays the superior role (*i.e.*, an artificial society without social structure). The former anticipation is essential for a game theoretical approach of institutional change as understood in old Veblenian institutional economics, *i.e.*, basically a co-evolutionary process between agency and structure leading to differentiation in the population of agents. Particularly, we model institutions as governed social structures with clearly codified entry and exit conditions for agents as members (compare Hodgson ([22], p. 18) for this particular aspect of institutions). To this extent they represent generic regulatory mechanisms that make societies stable on a large scale. Thereby, it is important to note that institutions are neither conceived as general-purpose vehicles, nor just as spontaneously emerging and exiting, but underlie individual life-cycles. Respectively, they evolve within a contingent path-dependent process that is dependent on the general level of societal trust. This aspect of accumulation makes agents endogenously heterogeneous and institutions diverse evolving aggregate structures, see [23] for a differentiation between heterogeneity and diversity.

In our model, agents populate an abstract topological space and interact with each other locally on a regular grid with linked edges (a torus) to avoid edge cell problems. The interaction is based on a prisoner's dilemma logic, *i.e.*, in every time step agents play the prisoner's dilemma game with their von-Neumann neighbors. According to the logic of the ordinary 2×2 prisoner's dilemma, agents can either cooperate or defect. In our model, agents are endowed with cognitive capabilities (a memory of events in the recent past, and a decision mechanism using this memory), which feed their individual decisions. In the course of the simulation, different agents accumulate different memories due to individual spatial interaction, and thus naturally evolve into a heterogeneous set of individual decision makers¹. Repeated cooperation between agents builds up trust, which in turn influences the emergence and exit of institutions as exclusive governed structures. It is important to distinguish between "institution-building proper", which by itself just constitutes part of the "rules of the game" [24], of the simulation, and its materialization as some special form of governance. Thereby, some members of this institution enforce compliance to the rule set. The special form, the realization of an institution, needs to be modeled explicitly by some agents taking over the role of enforcers, the role of executive power. As history teaches, executive power is needed for two distinct tasks: (1) It guarantees internal stability (compliance to the institutional rules); and (2) it warrants security from external threats (others trying to invade from outside). The institutional apparatus necessary to exert executive power is always financed by tribute payments of its members to their ruling executive. With a similar (and consistent) logic the model also takes care of the possibility of the break-up of institutions.

¹ Heterogeneity is thus not only an exogenous assumption replacing the (mainstream economic) assumption of a set of homogeneous representative agents; it indeed is a *process*, which evolves as part of the overall dynamics.

In this respect, we follow Hooper *et al.* [25], who highlight that leadership may get accepted in cooperative groups, if it crowds out free-riding or coordination errors. Moreover, the authors specify in their model that members may prefer to pay for the supervision instead of staying in an unsupervised group. Thereby, agents accept a hierarchical organization of social complexity, which is in line with our approach. Still, we do not implement public good games, as in [25], where mutual monitors punish defectors, but institutions with hegemonic leaders enforcing cooperation within clearly structural bounds, instead of more loose groups of agents sharing common properties. To this extent, our approach wants to emphasize the explicit character of a governed institution in comparison to more implicit group selection dynamics. However, our model shares some basic characteristics with [25], e.g., enforcement of cooperation becomes more costly with increasing group size, as we will outline in more detail. Where [25] investigates just one emerging group and the dynamics within group members and agents outside the group, we are able to investigate a whole number of institutions in a common spatial environment. This constraint is given by the methodology of dynamic evolutionary games in continuous time, building upon [26]. In the agent-based framework, time is mostly considered discrete and various interaction topologies can be implemented. Elsewhere, Smaldino and Lubell [27] show that a multi-agent approach may indeed help to investigate a diversity of solutions to social dilemmas. The authors investigate an “ecology of games”, where each game is analyzed with two different institutional mechanisms, capacity constraints and observation of behavior. The model we put forward implements the former institutional mechanism building upon a capacity constraint, *i.e.*, in our case a diminishing leader influence based on distance (“leadership distance decrement”) since we work with an interaction topology (2D space with finite, small neighborhoods) that clearly contrasts our approach from the aforementioned ones.

Our model is based on Sanchez-Pages and Straub [28], who analytically investigate the emergence of institutions in a multi-stage one-shot game where homogeneous agents are pairwise matched to play a game of prisoner’s dilemma. Each of the two agents participating in the prisoner’s dilemma (PD) has the choice between the two actions of cooperation (C) and defection (D). Since the game is played simultaneously and communication is prohibited, *a priori* the two players are not aware of their respective opponent’s choice of action, therefore, starting in a Hobbesian state of nature. If both players cooperate, they both achieve a payoff of R (reward), if they both choose to defect, they both end up with a payoff of P (punishment). Finally, if one agent cooperates and the other defects, then the cooperator gets a payoff of S (sucker’s payoff) and the defector receives T (temptation). Payoffs therefore satisfy $T > R > P > S$ and $2R > T + S$. Which strategy is chosen depends on the exogenously given level of trust within the society in the model of Sanchez-Pages and Straub [28]. As agents have the same level of trust, they always choose the same strategy, thus, only the two symmetrical outcomes of mutual cooperation (C,C) and mutual defection (D,D) are can be realized. Every agent in our model on the other hand has an individual trust level, which evolves over time as a result of her past experiences, thus, all four possible outcomes are considered. However, in their static model agents have the option to establish an institution that enforces cooperation between its members. To this end, they must choose a leader whom they can delegate the work of enforcing cooperation. The leader may not participate in the PD game but it may set a fee that all agents willing to join the institution have to pay to at least cover his opportunity costs. Games between members of

the institution always reach the cooperative outcome. Games between a member of the institution and an outsider, however, are not under institutional supervision and are treated like games in the state of nature. For convenience, Sanchez-Pages and Straub [28] label the former case (enforced cooperation) as “formal games” (*i.e.*, game partners comply with the formalities/rules of the institution) and the latter, as well as games between two institution-less agents, as “informal games”. With this basic setup, the authors of [28] go on to analyze equilibrium solutions on the number of agents within the institution, optimal fees and threats of secession. While their approach is instructive with respect to a number of issues, it considers only the case of one institution *versus* no institution in a one-shot static game. We argue in favor of a dynamic approach to catch the subtleties of the emergence, life, and exit of such coalitions between individual agents. Furthermore, we are able to study the evolution of a whole society of agents and institutions over time and to analyze these societies using a kind of institutional demography. These are all aspects that are impossible to derive from the static game described above.

Since an analytical model of such a dynamic complex adaptive version would hardly be tractable mathematically, we resort to the method of ABM. In the model, the artificial world is represented by a two-dimensional grid on which the agents can move around freely. Borders are wrapped around so that the matrix topographically corresponds to a torus. If an agent happens to meet other agents within her von-Neumann-neighborhood she plays a game of PD with each of them. If a cluster of at least three agents exists, these agents may decide to become sedentary, choose a leader and build an institution. Members of institutions are able to leave the institution each time step and the leader of an institution is allowed to set a new fee in each period. In what follows, all steps are presented in detail.

2.1. Initialization

At the start of a simulation run, n_A agents are distributed randomly across the grid. The random numbers are drawn from a pseudo-random number generator following a uniform distribution. Each agent is endowed with a memory of size κ . In this memory the agent cognitively stores the opponents’ choices of the last κ informal games. We define informal games as games played between (1) two agents who are not members of an institution; (2) an agent who is member of an institution and an agent who isn’t; or (3) two agents who are members of different institutions. In short, informal games are those games that are not supervised by the same leader. On the other hand, games played by two agents, who are members of the same institution, *i.e.*, those games where the cooperative outcome is enforced, are labeled formal games.

We further define the share of cooperative actions stored in an agent’s memory as her personal value for α , in this respect it is not comparable to the institutional mechanism of reputation suggested by [27], since we do not track the reputation of $n:n$ encounters. In contrast, we model the agent’s memory as personal perception of trust in the society based on past encounters independent from intersubjective reputation. If we assume, for instance, each agent to have a memory of the last ten informal encounters, *i.e.*, $\kappa = 10$, then $\alpha = 0.6$ is equivalent to the case that in any six out of the last ten informal encounters the agent’s opponents cooperated. The size of memory thus represents an assumption on the flexibility of an agent to adjust to new experiences. In this way, we endogenize

the evolution of trust α according to those new experiences of an agent. If, for instance, an agent meets a lot of other agents who cooperate, her personal α , *i.e.*, her trust in society, will rise and the agent will be more likely to cooperate in the future. Since we state that only informal games are memorized, we assume that enforced cooperation within an institution does not influence an agent's personal level of trust in cooperation between strangers. Obviously, at the initialization of a simulation run, no games have been played and therefore no actions would be stored in the agents' memories. We start with all agents having the same initial value of alpha at the start, which is a simulation parameter. With this starting value, we construct a random history of encounters for each agent, *i.e.*, a hypothetical history of events that corresponds to the given value of her personal α_0 at initialization. The histories are a list of encountered strategies, *i.e.*, the strategies that opponents played during the last κ informal encounters and are stored as a string with length κ , where each character represents a past encounter (with "C" representing an encounter with a cooperating agent and "D" having faced a defecting agent). Each agent has only one such list for all informal encounters, which contains no information to identify former opponents. Though the two extreme exemplary histories "CCCCCCDDDD" and "DDDDCCCCCC" both represent the same $\alpha_0 = 0.6$ (60% remembered cooperative informal encounters), they still represent heterogeneous histories. In the first example another cooperative encounter would push out the oldest remembered encounter (*i.e.*, the oldest memory would be forgotten) and change agent i 's memory to "CCCCCCDDDD" ($\alpha_i = 0.7$), whereas, in the latter case, it would change to "CDDDDCCCCCC" ($\alpha_i = 0.6$). Note that, although these encounters are stored in the proper historical order, the agents do not have any information as to how much simulation time has passed since the individual encounter happened, *i.e.*, the agents have an event based memory.

In contrast to the perfectly homogeneous agents in [28], the agents in our model are heterogeneous with regards to their location, their personal history and trust level within the simulated world. Though simulation time progresses in discrete intervals in our simulation, we employ full asynchronous updating with random ordering for our agents, which means that every agent performs all her actions in one go, before the next agent is activated. The consequences of all actions are effective immediately, *e.g.*, informal games have an influence on the memory, and, thus, α , of encountered agents, whether they have already been active in this round or not. This is a much more realistic assumption than synchronous updating, which introduces game phases for specific agent's actions during each round. The latter method was shown to be very problematic with some simulations by [29], who showed that some simulation outcomes could only be reached because the agents were activated in a specific static non-random order and were updated synchronously.

At the beginning of each time step, the activation order of every agent is shuffled randomly. Then, every agent is activated and takes all her actions before the next agent is activated. The following subsections describe the actions, which every agent may take during her round.

2.2. Movement

An agent who is currently not member of an institution, looks for an unoccupied site within her immediate von-Neumann-neighborhood (with a sight of one), randomly selects one of these and moves to its location if at least one such file exists. Thus, the agent can only move to a free location

directly above, below, right, or left of his current location or not at all if all cells are currently occupied.

2.3. *Playing the PD*

Leaving institutions aside for a while, the next step lets the agent play a game of PD against each of her von-Neumann-neighbors in random order, *i.e.*, up to a maximum of four games per round. In informal games, each agent plays a mixed strategy of cooperating with probability α and defecting with probability $(1 - \alpha)$. As stated above, the parameter α evolves endogenously for each agent. This setup stands in contrast to [28], who consider only cases of mutual cooperation or mutual defection, whereas our model allows for the cases of (D,C) and (C,D) as well.


Please note that, in contrast to [28], in our model, leaders are allowed to play PD games. The reason is that leaders can also have informal encounters and, thus, changes in their trust level (α_i) during their time as leaders of an institution. It would be counterintuitive that only the leaders are isolated from this societal evolutionary learning process.

2.4. *Building an Institution*

A cluster of at least three agents connected through their von-Neumann-neighborhoods may decide whether to build an institution. This cluster is formed by all agents directly or indirectly connected to each other, who currently are not members of an existing institution, *i.e.*, there is a path (unbroken chain) that traverses only members of the cluster, each of who is a von-Neumann neighbor of the former. An example for such a small cluster can be seen in Figure 1, which shows properly positioned neighbors (*i.e.*, members of the cluster) in green and unreachable neighbors in red.

An institution warrants enforced cooperation between its members at the cost of a membership fee. The process of institution formation proceeds in four steps. (1) Each agent within the cluster calculates if it pays to participate in the future institution; (2) Each agent willing to join the institution proposes a fee she would collect from the members of the institution, in case that the agent would become leader; (3) The agent proposing the lowest fee is appointed as leader; (4) Each agent aside from the leader decides whether to effectively participate in the institution under the designated leader and her proposed fee. If after these four steps, a connected set of two members and the leader remains, *i.e.*, at least three agents, then this connected set becomes a formal institution.

Figure 1. Example of a small cluster of von-Neumann-neighbors.



2.4.1. Step 1: Decision of Participation

At first each of the agents in the cluster calculates if it pays to participate in the future institution by comparing her potential informal payoff with her potential formal payoff as a member of the institution. The agent assumes that her further encounters will be similar to previous ones, *i.e.*, she expects to encounter the same mixed strategy that she herself currently employs. Thus the agent expects that she will encounter a cooperating agent with probability α and a defecting agent with probability $(1 - \alpha)$. She then sums up the four possible payoffs weighted by their expected probability to occur. This results in an expected potential profit of a single informal encounter of agent i with another agent as given in Equation (1).

$$\pi_i^I = \alpha_i[\alpha_i R + (1 - \alpha_i)S] + (1 - \alpha_i)[\alpha_i T + (1 - \alpha_i)P] \quad (1)$$


With R, S, T and P being the individual payoffs of the prisoners dilemma game for the respective situations. Superscript I stands for informal profit expectation as compared to superscript F for formal profit (within an institution).

With regards to formal games, we assume the quality of enforcement of cooperation to decrease with agent i 's distance $\delta_{i,L}$ to the leader L . The distance to the leader is measured as the shortest path between the agent and the leader that traverses only members of the institution, each of who is a von-Neumann neighbor of the former. The payoff for a formal game is then given in Equation (2).

$$\pi_i^F = \begin{cases} R(1 - \lambda\delta_{i,L}), & (\lambda\delta_{i,L} \leq 1) \\ 0, & (\lambda\delta_{i,L} > 1) \end{cases} \quad (2)$$

The parameter $\lambda \in [0,1]$ is exogenously given and serves as a weight for the loss in quality of enforcement, *i.e.*, a decrease of leadership effectivity with increasing distance (leadership distance decrement). Thus, a value of $\lambda = 1/3$ means that agents with a distance of at least three fields do not receive any payoff from formal games anymore, though, in the model at hand, it is only relevant at which distance the expected payoff of institutional cooperation drops below the expected payoff of informal games. Nevertheless, if only taking this decrease of effectivity into account, the above case would lead to a maximum von-Neumann distance of two fields from the leader, and a maximum institution size of 12 members and one leader. This can be seen on left hand side of Figure 2, which shows the von-Neumann distances from the leader in the center. It follows that smaller values of λ enable the emergence of much larger institutions. However, it is highly unlikely that institution size will ever come near to the theoretical maximum for that λ value as indicated before since the expected formal payoff is very low further away from the leader. Thus, the value of λ is more relevant for the question how much of the reward payoff remains on the cells near the leader, e.g., with a high value of λ even agents directly adjacent to the leader would gain very little or even too little payoff from formal games. The remaining share of the reward payoff, with respect to the distance from the leader $(1 - \lambda\delta_{i,L})$ —again for the example of $\lambda = 1/3$ —can be seen on the right hand side of Figure 2 ($2R/3$ if 1 field away, $R/3$ if two fields away, 0 if three fields away, with 0 not being displayed).

Figure 2. Exemplary von-Neumann distances and remaining payoff in the case of $\lambda = 1/3$.



Since the agent might have multiple encounters per round, she tries to anticipate the number of formal and informal games she will play actively and passively² in the next round. We assume that the agent takes all available local information into account when calculating her expected payoff. The expectation for the overall payoff in non-member state is then formulated in Equation (3).

$$\pi_i^{NMS} = \pi_i^I * n \quad (3)$$

With n representing the number of local neighbors, *i.e.*, the current number of agents occupying adjacent fields.

The expectation for overall payoff in the member state on the other hand is then given in Equation (4)

$$\pi_i^{MS} = ((n - m) * \pi_i^F) + (m * \pi_i^I) - \varphi \quad (4)$$

With φ being the fee that will be charged by the leader to-be and m representing the number of those local neighbors that will almost definitely lead to informal games—against agents that currently are in another, already existing institution. Thus, $(n - m)$ is the optimistically expected number of future co-members of the institution to-be, *i.e.*, the current number of non-members in their von-Neumann neighborhood.

Since the institution has not come into existence yet, and no leader has been chosen, the distance to the leader is not yet defined. Thus, for the first estimates of π_i^F and π_i^{MS} each agent assumes that $\delta_{i_L} = 0$, *i.e.*, that she might be appointed leader. This estimate will be corrected and decisions will be reevaluated as soon as a potential leader was selected, before the institution is formed, *i.e.*, wrong estimates will not have a huge effect.

The initial fee φ for this first evaluation of π_i^{MS} is again based on the assumption that the agent herself might be leader, and will also be proposed to the other agents in a later Step (2.4.2). We assume that an agent estimates her fee proposal according to the sum of collected fees equaling the cost of being the leader, *i.e.*, we assume that leaders do not factor in a profit margin. The cost of enforcing cooperation is then given in Equation (5).

$$c(\bar{\delta}, s) = \bar{\delta}\sqrt{s}, \quad (5)$$

² Passive play occurs when a game is initiated by another agent—*i.e.*, a game that is played outside of the target agent's turn.

Equation (5), thus, denotes the cost accruing to the leader from enforcing cooperation and s is the size of the institution, *i.e.*, the number of members including the leader. Furthermore, $\bar{\delta}$ represents the average distance of the leader (*i.e.*, herself) to all members of the institution. Obviously, the chosen cost function is just one of many possible alternatives, but it serves as a first reasonable and parsimonious approach. Future research may well investigate the effects of different cost functions.

The initial fee φ for this first evaluation of π_i^{MS} is then given in Equation (6).

$$\varphi = \frac{c(\bar{\delta}, s)}{s - 1} \quad (6)$$

We assume that the leader agent herself does not need to pay the fee, thus, the costs are divided by the number of members without the leader ($s - 1$). Since, at this moment, it is not yet clear to the agents how large the institution will, in fact, be (since potential members-to-be will be allowed not to join, once the potential leader is known), they use the size of the cluster they are part of as an estimate.

Each agent in the cluster evaluates the benefits of participating in the institution now and compares her expected non-member state payoff π_i^{NMS} with her expected member-state payoff π_i^{MS} . Only if the latter exceeds or equals the former, the agent is willing to participate in the institution. If all agents in the cluster come to the conclusion that they will be better off as members of the institution the formation of the institution proceeds to the next evaluation steps—else the institution will not be evaluated any further. Due to the completely asynchronous nature of the simulation, another agent might again start the evaluation of the same—or a slightly changed—cluster within the same time step.

2.4.2. Step 2: Proposing a Fee

During Step 2, each agent willing to participate in the institution proposes the fee that she would collect from members, which has already been calculated in the evaluation Step of 2.4.1., as given in Equation (6).

2.4.3. Step 3: Appointing a Leader

Next, the agent proposing the lowest fee is appointed as the leader of the potential institution to-be. If more than one leader proposes the lowest fee, one of them is appointed randomly.

2.4.4. Step 4: Final Evaluation

Finally, each agent in the cluster recalculates the distance to the leader to-be and reevaluates her expected payoffs. If the situation changed and her expected member-state payoff is no longer higher or equal to her expected non-member state payoff, she chooses not to join the potential institution. This procedure serves as a preliminary proxy for future implementations with more complex evaluation procedures, such as voting. Any agent that is disconnected/isolated from the cluster, due to other agents leaving the cluster, will also not be part of the institution to keep institutional territory

contiguous. Only if after this final evaluation a connected set of at least three agents (including the leader) remains, an institution emerges. All agents participating in an institution become sedentary and remain so until they eventually leave the institution or the latter breaks apart. Every agent may only leave, join, or found an institution once per round, which is evaluated in the given order.

2.5. Joining an Already Existing Institution

If an agent is located in the von-Neumann neighborhood of a member of an institution after moving, the former may choose to join the institution as well. Again, this agent compares her non-member-state payoff with her hypothetical payoff from joining the institution, based on actual distance to the leader and future size of the institution. If the member-state payoff is higher than or equal to the non-member-state payoff, the agent joins the institution.

2.6. Leaving an Institution (Re-Evaluating Membership)

In each time step, every member of an institution re-evaluates her gains from participation in the institution. If due to changed circumstances (e.g., changed neighborhood or individual trust), her member-state payoff no longer exceeds or at least equals her non-member-state payoff, the agent chooses to leave the institution. All members not connected to the leader anymore are forced to leave the institution as well. If the size of an institution falls beneath three, it ceases to exist.

2.7. Re-Evaluating the Fee

In every period, each leader of an institution re-evaluates the fee she collects from the members of the institution. If the sum of fees collected in the previous period is smaller than the cost accrued to it for enforcing cooperation, the leader would suffer a loss. In this case, she raises the fee such that the collected fees would equal the cost in the current period.

Since we employ strict asynchronous updating, each agent may actively initiate one attempt to found a new institution during his turn. If such an attempt results in the foundation of an institution during one agent's turn, it is founded effective immediately for all agents that finally participated. Agents who are activated after they passively participated in the founding of an institution can only play the PD during his turn, because agents may only participate in one institutional activity per time step, *i.e.*, founding, joining, leaving—evaluation alone is not considered an institutional activity for this purpose.

3. Simulation Experiments, Data Analysis and Results

In this section we illustrate the data analysis of simulation experiments and show the main results of the model. Simulations were performed with a payoff matrix for the pairwise prisoner dilemma that is in line with the non-degeneracy condition according to [1]³. The initial parameters for our

³ Basically resulting into $R - S - T + P \neq 0$ and $R - T - S + P \neq 0$, for details compare [1]. This condition has strong consequences for the calculation of the informal payoff of agents. If we would pick payoff values not meeting

experiments comprise settings given in Table 1. The memory size of agents (their cognitive endowment) smoothes the level of societal trust within simulation runs and works as a behavioral inertia therefore. Particularly a higher memory size leads to fewer fluctuations in societal trust, since agents play in a more consistent way. In this respect more agents are committed to a common average societal trust. Adjustments in low memory sizes have a strong effect, whereas adjustments in high memory sizes have weaker effects; e.g., changing the memory size from 80 to 100 encounters does have a weaker effect than changing it from 10 to 20 (especially since very low numbers of remembered encounters could theoretically even result in a completely new set of memories within a single simulation step). This result is quite intuitive, for that reason the memory size is not used as an additional variable for permutations in the following analysis.

Table 1. Initial parameters for simulation experiments.

World and Population		Prisoner Dilemma	
Worldsize (cells)	50×50	PD “punishment”	2
Population density	25%	PD “sucker’s payoff”	0
Agents	625	PD “reward”	4
Agent memory	20	PD “temptation”	5

The motivation of the analysis put forward is to get a deeper understanding of the generic logic of institution demography, the interdependent effects of individual strategic action, societal trust and institutional accumulation. The following two initial parameters are crucial for this investigation: the initial level of societal trust (α_0) and leadership distance decrement (λ). Correspondingly, we kept parameters given in Table 1 constant and computed simulation runs around couples of initial α_0 and λ with varying permutations, see Table 2.

As elaborated in Section 2, the initial level of alpha (α_0) fixes the starting distribution of individual cooperation and defection of agents. Otherwise the leadership distance decrement (λ) constrains the size of institutions during simulation runs. The two parameters serve as proxies for the dynamics of institutional life-cycles in the model. Experiments have indicated that the simulation may eventually reach a final state, either with static institutions (*i.e.*, complete cessation of institutional change) or without any institutions (only informal encounters). In order to analyze the dynamics in between these two states with more detail we could limit the parameter space due to experiences with former experiments and induce permutations of initial societal trust with a lower level of $\alpha_0 = 0.2$ for more stable institutional settings. These simulations, experiments 1–3 (Table 2), are varied with $\lambda = 0.15$; $\lambda = 0.2$ and $\lambda = 0.25$. Low leadership distance decrements result into a gain of power in the periphery of an institution (*i.e.*, a higher remaining profit in the periphery) and therefore lead to larger institutions. In such settings the simulation almost instantaneously tends to converge to a static institutional scenario. Higher λ leads to smaller institutions which are not stable over time, simply because the maintenance costs do not pay off anymore. This lead us to the expectation that simulation experiments 4–6 would deliver the most exciting results in terms of a

the “nondegeneracy” condition (e.g., $R - S - T + P = 0$) then the informal payoff function would be in linear dependence of alpha (trust), *i.e.*, a very special and singular case.

complex evolution (between order and chaos). The simulation experiments 7–9 have been initialized with $\alpha_0 = 0.4$ (that means a very high initial level of trust), where participation in an institution is not tempting.

Table 2. Simulation experiments with permutations of α_0 and λ .

Simulation Experiment	Initial Alpha (α_0)	Leadership Distance Decrement (λ)
1	0.2	0.15
2	0.2	0.20
3	0.2	0.25
4	0.3	0.15
5	0.3	0.20
6	0.3	0.25
7	0.4	0.15
8	0.4	0.20
9	0.4	0.25

We have looked into the evolution of some crucial variables with 20 replications (with randomly generated seeds) for each simulation experiment running over 10,000 rounds/periods, resulting in 9×20 datasets. Preliminary experimentation has shown that the institutional landscape may stabilize after approximately 5000 simulation periods, to this end we decided to gather an overview for the long run and have doubled this crucial time horizon. For every dataset we have created aggregated measures for the evolution of the following variables: mean societal trust, mean share of non-members in the population, mean number of institutions, mean age of institutions, and heterogeneity within the age structure of institutions. At first we are interested in the evolution of societal trust.


3.1. Societal Trust

Societal trust (the average share of cooperative actions within the whole population) may theoretically converge towards $\alpha = 1$ (full cooperation). Otherwise $\alpha = 0$ (full defection) never occurs, because agents will always take refuge in institutions once the amount of defectors increases in the population. The evolution of societal trust depends on the initial values of trust (α_0) and leadership distance decrement (λ). In general, the higher α_0 the higher is the level of societal trust during a run, since the initial memory of cooperation has a strong path-dependent effect and agents are not likely to build institutions respectively. Figure 3 shows the evolution of societal trust (α) for all nine simulation experiments, where (a) summarizes the permutations of λ for $\alpha_0 = 0.2$, (b) the permutations for $\alpha_0 = 0.3$ and finally (c) for $\alpha_0 = 0.4$. On the left hand side of Figure 3 we see graphs showing the evolution of the means over 20 replications and on the right hand side we see the standard deviations for these means accordingly. The data confirms that a higher initial value of societal trust leads also to higher values during the simulation when we compare Figure 3a–c. In a very similar manner a higher initial level of societal trust makes the evolution of trust more robust, as a comparison of the standard deviations in Figure 3a–c shows. However, the cause of this effect lies in the instability of institutions. Institutions are instable only if cooperation is high within the

population, because they are not tempting enough for agents to join (expected payoffs of non-member state exceed expected payoff of member state, as given, e.g., in the simulation experiments in Figure 3c. A high level of societal trust with a given memory size of $\kappa = 20$ encounters indicates a stable and robust process. Otherwise, Figure 3a shows, that if agents defect, more due to a low initial level of societal trust, they take refuge in institutions (enforcing them to cooperate) in early phases of the simulation.

Taking a closer look lets us follow that a lower λ leads in general to lower levels of societal trust within the simulation experiments. In Figure 3a,b the lowest level of societal trust is given in the cases of $\lambda = 0.15$, whereas $\lambda = 0.25$ leads to the highest levels of societal trust in all permutations. Of course this relation corresponds with the basic model setup as indicated in Section 3.2: lower λ results in larger institutions (smaller loss of power dependent on the distance between member and leader). Due to the spatial interaction topology (*i.e.*, locality effects), it is more likely that agents become members thereafter. In the very long run, high α_0 leads to processes where societal trust tends to converge to full cooperation.

Figure 3. Evolution of societal trust (α).



3.2. Evolution of the Population Structure

The next variable we investigate concerns the evolution of population structure, with regards to membership in institutions. Figure 4 shows the average share of non-members in the population across the conducted simulation experiments. It indicates that the population structure is quite stable over time, as the standard deviations show in Figure 4a–c. If leaders have greater influence in the periphery of institutions (low λ) then the amount of non-members playing with randomly encountered agents settles on a low level correspondingly.

Furthermore, we have checked for the robustness of this result by looking into the average number of institutions respectively, as illustrated in Figure 5. Since the number of institutions further stabilizes after an adaptation phase of 2000–3000 periods (also visible in Figure 4) we can conclude that there is a direct relationship between leadership distance decrement (λ) and the amount of institutions indeed as expected. Figure 5a–c show that permutations with $\lambda = 0.25$ result into 80 institutions, although this number deviates between 0–30 dependent on α_0 . Otherwise lower λ is more sensitive to the initial level of societal trust (α_0). This analysis further indicates that interesting cases of permutations are appearing around $\alpha_0 = 0.3$ and $\lambda = 0.2$. These cases have to get investigated in more detail by looking into the mean age of institutions.

Figure 4. Share of non-members in the population.






Figure 5. Number of institutions.

3.3. Institutional Life-Cycles

Since the population structure is rather stable, we have anticipated that the dynamics are more volatile in the evolution of institutions, dealing with entry, (in)stability and exit of institutions. Figure 6 shows the spectrum of institutional life cycles in a grand comparison of all computed permutations in simulation experiments 1–9. Accordingly, we have derived a deeper analysis on the mean age of institutions, in particular the accumulation process dependent on individual action and social structure. This full picture illustrates that our model produces two main results concerning the demography of institutions. On the one hand, institutions accumulate quickly from start due to a high number of free-riders in the initial population. Thereby, agents take refuge in governed institutional structures under leader supervision preventing exploitation in non-member state. Simulation experiments 1 and 2 result into this finite state of the simulation, where institutions do not fall apart anymore, *i.e.*, their life-time is infinite. We call this result a *static and ordered scenario of institutional change*. On the other hand, simulation experiments 8 and 9 indicate that institutions cannot stabilize, because their leaders are not able to influence the periphery of their institutions effectively, thereby, the sum of collected membership fees remains too small for long maintenance of


larger institutions, or ultimately any institution at all. In this case institutions have a very short life-time, they “pop up” very frequently in pulses. This final state of institutional change delivers the second main result in the “*in silico*” analysis of the model. We call this result a *dynamic but highly fluctuating scenario of institutional change*.

Figure 6. Mean age of institutions—Full picture long run.



However, we further identify a third result characterized by a more complex behavior. This result is best illustrated via the system behavior of simulation experiments 3, 5, 6, and 7. We call this result a *dynamic and complex scenario of institutional change*. Here, institutions cycle over a long period of time where it is undecidable to which state they may converge. Simulation experiments 5 and 7 converge to the highly fluctuating (almost chaotic) state of institutional change. We dub this process *institutional learning*, since the early advancing life-time of institutions experiences a turning point, where societal trust reaches significant levels that agents are tempted to get rid of social structures and switch to non-member status. Institutions in this model support social learning, which will be described in more detail at the end of this section. Otherwise, experiments 3 and 5 tend to develop towards the static and ordered scenario of institutional change, although with continuous fluctuations in the mean age of institutions. In order to confirm these results we provide an exploded view of this picture in Figure 7 and discuss some more details. Figure 7a–c highlight these results by showing just the medium run where the diversity within the mean institutional age becomes better visible.

Figure 7. Mean age of institutions—Exploded view medium run.



The most interesting cases of institutional learning are given in Figure 7b,c with $\lambda = 0.2$. The standard deviations indicate that these processes are not as deterministic as it seems in comparison to those with lower or higher leadership distance decrement. In this generic model of institutional change institutions may act as learning vehicles even with higher values of initial societal trust. The diversity within the age structure of institutions is high in these cases fluctuating between 10 and 100 periods. A closer view on the heterogeneity of institutional life-cycles among institutions within each time step (Figure 8) confirms this finding.

We observe that the higher the initial level of societal trust, the lower the heterogeneity within the age structure. Thus different values of λ may lead to phase-transitions towards another scenario. To conclude we want to turn the attention to the similarity in the trend of the aforementioned processes, which is not recognizable from the aggregated measures shown previously. Figure 9 gives the overall trend of replications within simulation experiment 5. It shows that institutions change indeed in cyclical processes, what we have dubbed *institutional learning*. The accumulation of these generic governed structures follows a similar trend, the cyclical behavior is of course leveled out in the aggregated views we have shown previously.

Figure 8. Mean age of institutions—Heterogeneity within the age structure of institutions.





Figure 9. Similarity in trend for $\alpha_0 = 0.3$ and $\lambda = 0.2$ over replications.



The diversity within institutional change that our model produces is a distinct feature that becomes evident only in a heterogeneous multi-agent configuration. The two main results indicate the deterministic corner solutions that can also be derived from a closed form equation-based game-theoretical model in continuous time, but an analysis of the complexity between these

deterministic solutions is out of range in this type of models. In this respect the discovery of *institutional learning* makes our study distinct in this realm of game-theoretical inspired evolutionary institutional economics and political economy. Again it is the interdependent interplay between agency and structure in space and time, which delivers such insights. The complex adaptive system dynamics in our model open a spectrum of potential institutional life-cycles over time. A crucial feature given by the model concerns the cognitive capabilities of agents, their potential to learn adaptively from the past and reevaluate their membership status in particular. Learning is, here, considered not just as a temporal processes, as mostly conceived in population games, but is severely dependent on the dynamically changing spatial distribution of agents, because institutions emerge and exit on certain places in the artificial political economy. Their bounded rationality in terms of [30] develops due to spatiotemporal adaptation (a mesoeconomic process), thereby depending on the complex evolution of the system as a whole but also on their institutional subsystems [31]. This finding is still in line with the socio-cybernetic theory of institutional change by Veblen, advanced by the means of evolutionary and complexity theory. Such a transdisciplinary approach to political economy may further stimulate novel modes of teaching in economics, see [32]. However it's these properties that are responsible for the likely volatile system dynamics, but they are also the major contributors to more noise and more complicated data analysis. Since the analysis of contingent path-dependent processes and emerging structures lies at the heart of evolutionary institutional economics and political economy, people are aware that results and moreover interpretations are never unambiguous.

However, we are able to conclude some major dependencies and dynamics in our model of institutional change. The evolution of societal trust is the major driving force behind the building up of institutions. If trust runs high on average, the need for institutions as we designed them decreases, because agents don't need executive protection. Consequently, this logic also works in the opposite direction, if trust runs low, agents demand institutions, and the absolute number of them will increase. Although this observation seems trivial it has some crucial ramifications dealing with the frequency of emergence and exit, *i.e.*, institutional change in cyclical patterns. These dynamics are majorly dependent on the initial values of exogenous variables, like the initial level of trust and the leadership distance decrement. As our experiments have shown, this parameter space determines the different paths and processes of institutional change. According to them we can identify three scenarios of institutional change generated by the computational simulation of our artificial political economy. Interestingly, our three scenarios share some generic characteristics of the results given by the complex system analysis in [33]. Stuart Kauffman has shown that in ordered regimes the elements freeze very fast and form a bigger cluster, which spans across the system. In the chaotic regime there is no frozen component; instead, a connected cluster of unfrozen elements appears. Small changes in the initial parameters may lead to strong reactions of the whole system. Transitions from the ordered to the chaotic state are possible through phase transitions, where the transition region is called a complex regime. In this regime frozen and unfrozen elements are percolating simultaneously with very sensitive conditions on the complex edge between chaos and order. The three scenarios found in our experiments are correspondingly called:

- *static and ordered scenario of institutional change*
 → a non-cooperative world indicated by simulation experiments 1 and 2

- *dynamic but highly fluctuating scenario of institutional change*
 - ➔ a cooperative world indicated by simulation experiments 8 and 9
- *dynamic and complex scenario of institutional change*
 - ➔ complex institutional learning indicated by simulation experiments 3, 5, 6, and 7

To give a more vivid and intuitive understanding of the different cases compare Figures 10–13, which visualize the artificial landscape of exemplary runs of experiments 1, 3, and 9. The subfigures show snapshots of the landscape of agents and institution in different points of time—please note that there are different timespans between subfigures in Figures 10–13, since the different classes have gravely varying convergence speeds. The subfigures show non-member agents, who have not yet been member of any institution depicted as blue squares and non-member agents who have already been in an institution before depicted as red squares. Furthermore, green circles depict leaders and green squares depict members of institutions, with black connections between members of the same institution. The darkness of institutional leaders and members indicates the age of the institution—*i.e.*, light green reflects a young institution, where dark green depicts a relatively old institution. Figures 11 and 13 additionally show pink (light red) squares and circles, which indicate agents in attempted institutions, *i.e.*, institutions that only existed for one evaluation period (technically these have an institution age of 0). These come into existence when the overall level of α is very high and no stable institutions exist anymore and are disbanded immediately. When taking a look at the following figures, please keep in mind that we interpret the grid as a torus; it might look like there are one or two-person institutions at the edges of the landscape, but in truth they continue on the opposing edge of the landscape.

Figure 10, now, is an the visualization of the landscape of an exemplary run of experiment 1 ($\lambda = 0.15$, $\alpha_0 = 0.2$, other parameters as given in Table 1) and shows what we called a *static and ordered scenario of institutional change*. The first subpicture on the left hand side shows that after a very short time (at $t = 10$) most agents already are members of institutions, with a number of agents still wandering around. The next subpictures show that very quickly (at $t = 50$) most of the agents are now members of institutions, that most institutions are already relatively old and stable, with only very few agents still wandering around. Shortly after (at $t = 250$) all agents are members of institutions, which are completely stable (old). The society is now ordered, in the sense that there is only cooperation enforced in institutions and static since there are no elements of change anymore, *i.e.*, the landscape is identical for all following time steps. Since this lock-in happened so fast, there was almost no change in the societal trust level—a graph of the development of societal trust over time would show a flat line.

Figure 10. Institutional landscape of an exemplary simulation run ($\lambda = 0.15$, $\alpha_0 = 0.2$).




Figure 11 now shows the visualization of the landscape of an exemplary run of experiment 9 ($\lambda = 0.25$, $\alpha_0 = 0.4$, other parameters as given in Table 1) and shows what we called a *dynamic but highly fluctuating scenario of institutional change*. The first subpicture on the left hand side shows that after a very short time (at $t = 10$) most agents have already been members in very shortly lived institutions, with only a very small number of institutions (in this example only 1), because the initial level of societal trust was very high to begin with. The next subpicture shows that in the short run a small number of institutions emerged ($t = 250$). This was the result of an already highly heterogeneous distribution of α , which lead to agents with lower α to still seek shelter in institutions—during this process, societal trust even went down a little. Nevertheless, the system quickly returns to converge against full cooperation, while institutions have not needed anymore for a long time due to the already high level of trust.

Figure 11. Institutional landscape of an exemplary simulation run ($\lambda = 0.25$, $\alpha_0 = 0.4$).




Figure 12 shows the development of societal trust over time of the same exemplary simulation run, which was depicted in Figure 11. As can be seen, societal trust rises rather quickly, while the process of reaching full cooperation can take quite long in singular runs. There is almost no difference

between member and non-member agents in terms of average trust, since there are only few, short-lived institutions early on, *i.e.*, agents are non-members almost all of the time.

Figure 12. Societal trust timeline of an exemplary simulation run ($\lambda = 0.25, \alpha_0 = 0.4$).



Figure 13 now shows the visualization of the landscape of an exemplary run of experiment 3 ($\lambda = 0.25, \alpha_0 = 0.2$, other parameters as given in Table 1) and shows what we called a *dynamic and complex scenario of institutional change*. The first subpicture on the left hand side shows that after a very short time (at $t = 10$) a number of institutions already exist, with the majority of agents not being members of institutions. The next subpicture ($t = 250$) shows that a number of institutions have existed for some time, while there also are some quite young institutions, as well as a number of non-member agents. The following subpictures show that the number of institutions rose, since in this particular simulation run societal trust went down to quite low levels (at $t = 2500$). Nevertheless, in the medium run, as depicted in the following subpictures ($t = 10,000$ to $t = 15,000$), overall trust rose again, thus requiring less institutions and in the long run (at $t = 35,000$) the artificial society arrived at full cooperation again.

Figure 13. Institutional landscape of an exemplary simulation run ($\lambda = 0.25, \alpha_0 = 0.2$).






Figure 13. *Cont.*

Figure 14 now shows the development of societal trust over time of the same exemplary simulation run, as depicted in Figure 13. As can clearly be seen initially α increases and decreases again, which is then followed by a kick-in of the self-reinforcing process. In the long run, societal trust rises until full cooperation is reached.

Figure 14. Societal trust timeline of the exemplary simulation run ($\lambda = 0.25$, $\alpha_0 = 0.2$).



In summary institutions in our model indirectly support social learning. Those agents with a low α are more likely to join institutions and thus become sedentary—compare Figure 14, where it can be seen that the average α of non-member is consistently higher than that of members (in the timeframe in which institutions are needed). In case of larger institutions a number of these low α agents are even out of reach of non-members or members of other institutions, as the former are embedded deep within the institution’s realm. Thus, these agents on average play fewer or even no informal games at all, during their membership in the institution. Non-member agents with a higher α stay mobile and then have a higher chance to encounter other agents, which also have a higher α . Thus they learn to cooperate even more over time, eventually increasing overall trust over time. These agents then function as “emissaries” of cooperation, slowly increasing the trust of agents situated on the outskirts of institutions (“teaching” them to cooperate). Since these are further away from their leader, their expected formal payoff can be quite low—depending on the λ parameter, which leads

to larger institutions, but lower formal payoff. Thus, they can be quite sensitive to changes in cooperative behavior shown by informal games with wandering non-member agents and start to leave institutions quickly.

This notion seems to be quite intuitive since in the early phases agents tend to take refuge in institutions, because trust is very low, then trust rises and then the system experiences a crucial transition, because life outside of the institution promises higher short-term payoffs on average and agents start to cooperate even more, but without institutions. Since a personal $\alpha > 0.5$ means that the agent will cooperate more often than defect, and thus rather cause “positive” than “negative” memories in other agents. “Positive” memories lead to more cooperative behavior in other agents, which reinforces itself and, in the very long run, defection is crowded out and the society arrives at full cooperation, which is then stable.

On the other hand this scenario may also converge to a static state of frozen institutions quite quickly. This is the case when cooperation without institutions takes too high costs in the long run, or when defection is still beneficial. It is also likely to occur when setting the memory parameter (κ) to a small value, with the result that the wandering agents (what we called “emissaries” before) learn non-cooperative behavior through playing with members of institutions, while the level of α is still low. This can also lower the emissary’s α to such a degree that he himself wishes to join the institution, because he now feels “safer” as a member of the institution. Consequently societal trust may converge to a constant level, once all agents have founded or attached to an institution. However the dynamic attractors for a transition in the one or the other direction are sensitively depending on certain permutations of initial societal trust and leader influence in the periphery.

4. Conclusions

In our paper we have presented an agent-based model connecting the theory of games with institutional economics via socioeconomic computational simulations, thereby providing new insights for evolutionary institutional economics and political economy. We follow the core concept of the old institutional economics, where institutions are conceived as real social structures. The Veblenian idea of institutional change via cumulative causation of habits of thought serves as the dominant proxy for this emerging field of socioeconomic research. Obviously a synthetic translation of this concept into a concrete model and its subsequent simulation rigorously cuts components of its original semantic narrative. However, our findings indicate that even an approximate, at best sufficient, simulation of this kind of institutional economics demands a transdisciplinary approach by a formal combination of evolutionary game theory and complex systems theory. Such an endeavor can be realized via the use of computational methods and simulations of heterogeneous multi-agent systems in discrete time. A modern conception of complexity as composed by Wimsatt [34] may serve as a theoretical blueprint respectively, because he advances the Simonian [31] notion of reductionist heuristics as means to more holistic theories, in contrast to linear additive methods of total aggregation he refers to decomposition. This basic argument favors a theory of practice approach, which is again in line with the old institutional economics that builds to a great extent on the philosophy of American pragmatism.

Institutions appear as diverse stabilized social structures dependent on the individual evaluation of heterogeneous agents to enter or exit them. Then the iterated PD builds the formal proxy for sequential social conflict. Traditionally it is assumed that cooperation may crowd out defection in the long run under certain conditions, such as group selection or the existence of certain beneficial topologies. Still, from a political economy perspective it is very likely that continuous defection gets never crowded out under moral assumptions such as altruistic punishment [35] or indirect reciprocity [36]. On the contrary it is more likely that individual agents will build institutions (conceived as meta-stabilized structures of culturally transmitted rules) capable of supporting social learning, especially for large scale; *i.e.*, political economy. If such a structure is established once, it is highly realistic that there are costs associated to the maintenance of it. We regard this trade-off as essential for the dynamics of institutional change since agents continuously need to evaluate their participation in stable social structures for reasons of conservation or even progress (inertia *vs.* change as discussed extensively in [16]). To this extent our third scenario represents a very realistic case of institutional change, where institutions emerge and exit over long periods in a cyclical movement. These institutions support social learning, where in times of defection agents can take refuge in protected social structures and leave them in time of cooperation. Thereby, the role of a leader executing, such a policy via institutional structures may become more and more redundant once free-riders have “learned” the benefits of cooperation. Still, it is not said that even in such a scenario the system switches again towards a protective regime eventually. Nevertheless, we have to consider that this long run decline of institutional significance would very likely not hold in a model with agent mortality and procreation, where offspring would not be exact copies of their parents.

Due to the continuous spatiotemporal contact of individual agents with institutional structures, agents share and anticipate the current societal level of trust and learn from it, thus influencing their future behavior. Of course this is only possible if the agents are equipped with a simple cognitive apparatus and a crude memory of past events, at least. The framework presented and the simulations carried out clearly encourage further research in the envisaged direction. Most of the issues explicitly studied so far are just one example of what could be seen as essential features of emergence and exit of institutions in human societies. In our framework we use very basic concepts of power and exploitation regarding the institutional leader. Even in such an environment, where coercive power relations between agents are highly stylized—centralized and monopolized by a given ruling political entity—it is possible to observe a wide variety of forms of emergence and exit of institutions, *i.e.*, *diversity in institutional life-cycles*. Those basic political economy elements point out the potential of such an evolutionary institutional economic approach for further models of power and exploitation. The rules for the fee paid to the agent leading an institution as well as the rule for choosing this leader are just typical economic concepts to keep dynamics still simple for useful exploration. Extending our approach by replacing them with ideas reflecting pre-existing power relations securing exploitation would certainly enable our example to be a more adequate image of the historical emergence of institutions—where there has been a historical primacy of coercive power and exploitation compared to voluntary interactions on markets, thereby hinting at the potential scope of the approach. Extensions of this kind clearly suggest opening up new dimensions of heterogeneity of agents leading to a richer diversity of institutions.

Acknowledgments

The authors would like to highlight that initial research and previous versions of this paper were produced during a research project granted by the Oesterreichische Nationalbank (OeNB—the central bank of the Republic of Austria) within the “Jubiläumsfonds” project 13350. Furthermore we want to thank participants and organizers of the research area sessions of RA-C “Institutional Change” at the annual conferences of the European Association for Evolutionary Political Economy (EAEPE) during the years 2009–2012; as well as at the participants and organizers of the annual conference of the Association of Heterodox Economics 2009 for fruitful and productive comments. We further want to acknowledge the highly constructive critique of three anonymous referees.

Finally the authors are very grateful and acknowledge the efforts of Hardy Hanappi inspiring this research.

Author Contributions

The author’s contribution is equal in all aspects of the reported research and writing of the paper.

Conflicts of Interest

The authors declare no conflict of interest.

References

1. Young, H.P. *Individual Strategy and Social Structure*; Princeton University Press: Princeton, NJ, USA, 1998.
2. Gintis, H. *The Bounds of Reason. Game Theory and the Unification of the Behavioral Sciences*; Princeton University Press: Princeton, NJ, USA, 2009.
3. Maynard-Smith, J. *Evolution and the Theory of Games*; Cambridge University Press: New York, NY, USA, 1982.
4. Axelrod, R.M. *The Evolution of Cooperation*; Basic Books: New York, NY, USA, 1984.
5. Nowak, M.A. *Evolutionary Dynamics. Exploring the Equations of Life*; The Belknap Press of Harvard University Press: Cambridge, MA, USA, 2006.
6. Bowles, S.; Gintis, H. *A Cooperative Species. Human Reciprocity and Its Evolution*; Princeton University Press: Princeton, NJ, USA, 2011.
7. Richerson, P.; Henrich, J. Tribal social instincts and the cultural evolution of institutions to solve collective action problems. *Cliodynamics* **2012**, *3*, 38–80.
8. Ostrom, E. *Understanding Institutional Diversity*; Princeton University Press: Princeton, NJ, USA, 2005.
9. Hodgson, G.M. What is the essence of institutional economics? *J. Econ. Issues* **2000**, *34*, 317–329.
10. Dopfer, K.; Potts, J. *The General Theory of Economic Evolution*; Routledge: Abingdon, UK, 2008.

11. Hanappi, H. The neumann-morgenstern project—Game theory as a formal language for the social sciences. In *Game Theory Relunched*; Hanappi, H., Ed.; InTech: New York, NY, USA, **2013**, pp. 3–26.
12. Bourdieu, P. *Praktische Vernunft: Zur Theorie des Handelns*; Suhrkamp: Frankfurt, Germany, 1982.
13. Veblen, T. *The Theory of the Leisure Class*; MacMillan: New York, NY, USA, 1899.
14. Veblen, T. Why is economics not an evolutionary science? *Q. J. Econ.* **1898**, *12*, 373–397.
15. Dequech, D. The Demarcation between the “Old” and the “New” institutional economics: Recent complications. *J. Econ. Issues* **2002**, *36*, 565–572.
16. Heinrich, T.; Schwardt, H. Institutional inertia and institutional change in an expanding normal-form game. *Games* **2013**, *4*, 398–425.
17. Elsner, W.; Heinrich, T. A simple theory of meso: Co-evolution of institutions and platform size. *J. Soc-Econ.* **2009**, *38*, 843–858.
18. Dopfer, K.; Foster, J.; Potts, J. Micro-Meso-Macro. *J. Evolut. Econ.* **2004**, *14*, 263–279.
19. Arthur, W.B. Complexity economics: A new framework for economic thought. In *Complexity Economics*; Arthur, W.B., Ed.; Oxford University Press: Oxford, UK, 2014.
20. Wäckerle, M. *The Foundations of Evolutionary Institutional Economics. Generic Institutionalism*; Routledge: Abingdon, UK, 2014.
21. Aoki, M. *Toward a Comparative Institutional Analysis*; MIT Press: Cambridge, MA, USA, 2001.
22. Hodgson, G.M. What are institutions? *J. Econ. Issues* **2006**, *40*, 1–25.
23. D’Ippoliti, C. *Economics and Diversity*; Routledge: London, UK, 2011.
24. North, D.C. *Institutions, Institutional Change and Economic Performance*; Cambridge University Press: Cambridge, MA, USA, 1990.
25. Hooper, P.L.; Kaplan, H.S.; Boone, J.L. A theory of leadership in human cooperative groups. *J. Theor. Biol.* **2010**, *265*, 633–646.
26. McElreath, R.; Boyd, R. *Mathematical Models of Social Evolution: A Guide for the Perplexed*; University of Chicago Press: Chicago, IL, USA, 2007.
27. Smaldino, P.E.; Lubell, M. Institutions and cooperation in an ecology of games. *Artif. Life* **2014**, *20*, 207–221.
28. Sanchez-Pages, S.; Straub, S. The Emergence of Institutions. *B.E. J. Econ. Anal. Policy* **2010**, *10*, 1–36.
29. Huberman, B.A.; Glance, N.S. Evolutionary games and computer simulations. *Proc. Natl. Acad. Sci. USA* **1993**, *90*, 7716–7718.
30. Simon, H. Bounded rationality and organizational learning. *Organ. Sci.* **1991**, *2*, 125–134.
31. Simon, H. The architecture of complexity. *Proc. Am. Philos. Soc.* **1962**, *106*, 467–482.
32. Elsner, W.; Heinrich, T.; Schwardt, H. *The Microeconomics of Complex Economies. Evolutionary, Institutional, Neoclassical, and Complexity Perspectives*; Elsevier/Academic Press: Oxford, UK, 2014.
33. Kauffman, S. *The Origins of Order: Self-Organization and Selection in Evolution*; Oxford University Press: Oxford, UK, 1993.

34. Wimsatt, W.C. Aggregate, composed, and evolved systems: Reductionist heuristics as means to more holistic theories. *Biol. Philos.* **2006**, *21*, 667–702.
35. Fehr, E.; Gächter, S. Altruistic punishment in humans. *Nature* **2002**, *415*, 137–140.
36. Nowak, M.A.; Sigmund, K. Evolution of indirect reciprocity. *Nature* **2005**, *437*, 1291–1298.

Of Coordinators and Dictators: A Public Goods Experiment

Jürgen Fleiß and Stefan Palan

Abstract: We experimentally investigate whether human subjects are willing to give up individual freedom in return for the benefits of improved coordination. We conduct a modified iterated public goods game in which subjects in each period first decide which of two groups to join. One group employs a voluntary contribution mechanism, the other group an allocator contribution mechanism. The setup of the allocator mechanism differs between two treatments. In the coordinator treatment, the randomly selected allocator can set a uniform contribution for all group members, including herself. In the dictator treatment, the allocator can choose different contributions for herself and all other group members. We find that subjects willingly submit to authority in both treatments, even when competing with a voluntary contribution mechanism. The allocator groups achieve high contribution levels in both treatments.

Reprinted from Special Issue: Aspects of Game Theory and Institutional Economics, *Games*. Cite as: Fleiß, J.; Palan, S. Of Coordinators and Dictators: A Public Goods Experiment. *Games* **2013**, *4*, 584–607.

In 2005, a special issue of *Science* listed the 25 areas where scientists perceived the most important gaps in our knowledge to date [1]. These included the question, raised by Pennisi [2], of how cooperative behavior evolved to form the basis for the complex societal structures we observe today. She pointed out the importance of investigating which conditions and institutional settings promote cooperation in situations where individuals have an incentive not to cooperate. A famous example of such a dilemma is of course the contribution to a public good. In the standard setting, individuals have strong incentives to maximize their own payoffs by free riding and not contributing to the public good. As a result, a group of rational actors would be unable to supply a public good.

A large number of laboratory experiments have investigated cooperation in the public goods game (for reviews, see [3,4]). In the most common version of the repeatedly played public goods game, each individual in a group makes his or her own decision about how much of the endowment to contribute to a public good in every period. The results show that contributions tend to start out at an average of around 50% and decline towards zero [4,5]. Looking at individual behavior, a number of subjects are usually found to contribute in the first few periods of repeated public goods games. Over time, their contributions decline as they observe other subjects free riding and contributing nothing. In the end, because of these conditional cooperators' reactions to the free riders, the public good no longer gets produced [6–8].

These somewhat disappointing findings on human cooperative behavior in such dilemmas have been qualified by more recent results. There are mechanisms that can foster contributions to the public good. One such solution is monetary punishment, as introduced by Ostrom, Gardner and Walker [9]. Their paper, and a number of follow-up studies, show that (centralized and

decentralized) punishment and reward can stabilize contributions at high levels [10–13].¹ Besides punishment, voting on the implementation of different proposed contribution rules has been shown to have a positive effect [14,15].

We claim that besides the instruments of punishment and reward, direct power over the decisions of others can play an important role when it comes to the success of collective action in dilemma situations. Weber [16] defines power as “the likelihood that one person in a social relationship will be able, even despite resistance, to carry out his own will.” Structures of (asymmetric) power distributions are omnipresent in everyday life and characterize whole societies, but also groups, (business) organizations and the like [16,17]. Yet, despite its obvious importance in everyday life, the discipline of economics has not devoted much time to studying power over the decisions of others (for an analogous argument and another recent experimental study regarding power, see [18]).

Kroll *et al.* [19] show that contributions to a public good increase if group members in a public goods game can vote for a binding proposal as opposed to making voluntary contributions. One particularly noteworthy recent exception builds on the idea of binding subjects to specific actions. It employs a new contribution mechanism in public good games, based on an asymmetric distribution of power: the allocator mechanism. The two studies introducing this topic are Hamman *et al.* [20] and Bolle and Vogel [21]. Both show that, under certain conditions, one way of promoting the provision of a public good is to establish an allocator who has absolute power over the decisions of all group members. In the unique rational expectations equilibrium, this allocator is then able to force all group members to contribute their full endowment to the public good, thereby maximizing the collective outcome. Hamman *et al.* [20] and Bolle and Vogel [21] largely confirm this theoretical prediction and show experimentally that the use of an allocator results in comparatively very high contributions to the public good.

Where the two studies differ is in the specifics of group members’ and allocators’ choice sets and in the structure of the experiment. Hamman *et al.* [20] let group members elect an allocator and find that groups ensure full provision of the public good primarily by electing pro-social allocators. Since each group of nine holds a new election every period, their setting allows for punishment by removing underperforming allocators from power. Allocators who contribute fully for everyone are found to be re-elected in almost all cases. Bolle and Vogel [21] choose a different first phase for their experiment. They initially let subjects play 10 periods of a public goods game with voluntary contributions. This is followed by one period, where an allocator is chosen (either randomly or by election) to make the allocation decision for the two other members of her three-person group. This sequence of voluntary (10 periods) and allocator contribution phases (one period) is repeated twice, such that subjects play three allocator periods in total. Like Hamman *et al.* [20], Bolle and Vogel [21] observe higher contributions in the allocator setting than in the setting with voluntary contributions. Interestingly, they find no statistically significant differences between the election and the random selection treatments.

The great success of the allocator mechanism documented in these two studies merits further research. We explore its performance characteristics by (i) systematically varying the action space of

¹ For a recent review, see [4].

the allocator and by (ii) studying whether subjects prefer groups governed by the allocator mechanism over groups where they can freely choose their own contribution when group choice is endogenous. Note that the two precursor studies implement the allocator mechanism in a way that either forces all subjects to participate or allows endogenous participation, such that non-participating subjects profit from the public good of participating subjects. In this second case, Hamman *et al.* [20] find that the allocator mechanism is not able to increase contributions due to free riding. Only when communication is possible do half of their groups choose to transfer their decision rights and achieve high contribution levels. We build on these results and investigate whether the transformation of the public good into a club good, from which only group members can profit, is also able to foster the allocator mechanism's efficiency. Our second question therefore is of special importance, since it captures a subject's willingness to submit to authority for her own benefit and the benefit of the whole group. This question is also closely related to a major finding in the discipline of new institutional economics [22]. It states that the voluntary participation of subjects in finding a solution to coordination problems substantially increases the likelihood of success.

Such endogenous institution choice has previously also been examined for the punishment and reward mechanisms mentioned above. One approach is to let subjects vote whether they want to implement, e.g., punishment in the public goods game they will later be playing [23]. Another is to let subjects self-select into groups with different, exogenously fixed institutional settings. Gürerik *et al.* [24] find that subjects are more likely to self-select into groups with sanctioning institutions than into alternative groups and that the likelihood of choosing the group with sanctioning institution increases over time. In this way, they show that when two groups with different institutional settings compete against each other, the group with a sanctioning institution—due to the higher payoffs it generates for its members—prevails in the end. Hamman *et al.* [20] also present some of these aspects in their experiments. They allow subjects to choose whether they want to be part of electing an allocator who will then make the contribution decision on their behalf or whether they want to choose their level of contribution themselves. The important difference to the design of Gürerik *et al.* [24] (and this study) is that subjects who choose not to be part of the electoral delegation mechanism in Hamman *et al.* [20] nonetheless profit from the public goods contributions made by subjects who have delegated their decision power. This allows subjects who have not joined the delegation mechanism to free ride on its outcomes.² Without communication, Hamman *et al.* [20] obtain an average contribution level of only 11%. In this setting, the allocator mechanism thus fails to sustain high public goods contributions. Whether groups governed by the allocator mechanism have an advantage over groups with voluntary contribution when one group does not profit from the contributions of the other is an important and unanswered question.

Building on Hamman *et al.* [20] and Bolle and Vogel [21], we thus identify two important questions. First, do subjects prefer a group governed by the allocator mechanism over a group with a voluntary contribution mechanism? Second, which factors influence subjects' group choice?

² Kosfeld *et al.* [25] showed that centralized punishment institutions are able to prevail in a setting where non-participating players can free ride on the contributions of the players participating in the centralized punishment institution. This leads to high levels of efficiency in their public goods experiments.

The present article answers both of these questions. As an additional innovation, we drill down into the role played by the allocator's action space. Specifically, we compare a treatment with what we term a coordinator—an allocator who can choose one uniform contribution level for all members of her group, including herself—to a treatment with a dictator—an allocator who can choose a contribution level for herself and a different, uniform contribution level for all other group members. This mimics many settings outside the lab where a group leader or government establishes policies that apply to all group members equally (e.g., regulations that require every able-bodied male adult to contribute to the public good of national defense by having to serve a term in the military, as exists in many countries). Consider as a loosely-related example for our setting a country's decision to join the European Union. This country faces a tradeoff between giving up the freedom to decide on its laws and regulations entirely on its own and abdicating some of its regulatory authority to the EU institutions in return for the benefits from greater cooperation. This example shares with our design the feature that leadership of the group, in this case, the presidency of the council of the EU, rotates through all member states, with each country serving only one term. (Another example would be the decision by a stone-age human to join a tribe, thus giving up individual freedom in order to gain the ability to jointly hunt larger game, which is argued to have contributed to the rise of modern civilization; see, e.g., [26]).

The remainder of this paper is structured as follows. In Section 1, we state our research questions and derive our hypotheses. Section 2 outlines the experimental design and procedures. Results are presented in Section 3 and discussed in Section 4.

1. Research Question and Hypotheses

We investigate the question of how societal coordination can arise endogenously in response to economic coordination problems. We take a standard public good game as our workhorse model and augment it by giving subjects the freedom to select into one of two groups at the beginning of every period. In the voluntary contribution group (VCG), they play a standard public good game by deciding how much of their endowment to keep for themselves and how much to invest into a public good. If subjects select into the allocator contribution group (ACG), one group member is randomly chosen to set the contribution level for all ACG members.

Given a contribution level, we use the same payoff function in both groups. Specifically, a subject's payoff for any one period in our experiment is calculated as follows:³


$$\pi_i = E_i - c_i + \frac{\lambda}{n_\theta} \cdot \sum_{j=1}^{n_\theta} c_j^\theta \quad (7)$$

where π_i is the payoff of subject i , $E_i = E = 20$ is a subject's endowment in each period in experimental currency units (ECU), c_i is the subject's contribution to the public good in this period, $\lambda = 1.6$ is a constant determining the marginal per capita return (MPCR), n_θ is the number of subjects in group $\theta \in \{VCG, ACG\}$ and $\sum_{j=1}^{n_\theta} c_j^\theta$ is the sum of all contributions of subjects j in group θ in this period.

³ We suppress the period index in order to streamline the notation.

The return from the public good is rendered independent of the group size through the inclusion of n_θ in the denominator of the MPCR. It thus depends only on the average contribution in the group (this follows the design of Rockenbach and Milinski [27]). In the special case that only a single subject selects into one of the groups, the subject's contribution is automatically set to zero, and no public good is generated (subjects are so informed in the instructions). Note that subjects receive information about their group's size before making their contribution decision.

Figure 1. Treatment design and terminology.



Hamman *et al.* [20] and Bolle and Vogel [21] implemented the allocator decision in a way that allowed the allocator to set a different contribution for herself than for the other group members. This allows for the rise of “corruption”, which is how we refer to the case where the allocator does not contribute to the public good.⁴ Our design expands on this idea by modeling two different types of allocator decision options. We will continue to use “allocator” and ACG as the general terms, but will distinguish between a “coordinator” and a “dictator” treatment in our design. In the former, the coordinator can choose a contribution level, which then applies to all group members, including herself. In the latter, the dictator can choose two contribution levels, one of which applies to all group members excluding herself, while the other applies only to herself. In keeping with the vocabulary just laid out, we will be speaking of two forms of ACGs—the coordinator contribution group (CCG) and the dictator contribution group (DCG). Figure 1 illustrates this terminology. Finally, we wish to explore the impact of subjects' social preferences on their behavior in our experiment. For this reason, we elicit their social value orientation using the social value orientation (SVO) slider measure developed by Murphy *et al.* [28].

⁴ Note that our instructions generally contained neutral wording, for example, referring to the VCG (ACG) as the “group with individual contribution choice” (“group with contribution choice by a randomly determined player”).

1.1. Rational Expectations Predictions

To predict the group choice, we have to take a look at the expected contributions and payoffs in each of the two groups. In the VCG, the rationally⁵ expected behavior is not to contribute, yielding an expected payoff to every subject equal to her endowment. (This is also the minimax payoff in the VCG.) In the ACGs, there are different predictions for our two treatments. Given that the coordinator can only set one uniform contribution level for all group members, it is immediately apparent that for any $\lambda > 1$ (and assuming $E > 0$), the profit-maximizing strategy is to set the contribution level equal to the common endowment, E . The payoff to both the coordinator and the other group members, then, is the payoff from full cooperation:

$$\pi_{i,CCG} = \frac{\lambda}{n_{CCG}} \cdot E \cdot n_{CCG} \quad (8)$$

Given that $\lambda = 1.6$ and $E = 20$ in our setting, we thus derive the first part of our first hypothesis:

Hypothesis 1a. *In the CCG, coordinators always contribute the full endowment.*

In the coordinator treatment, the expected payoff as a member of the allocator group is higher⁶ than the minimax payoff in the voluntary contribution group, which equals the endowment E .⁷ Under rational expectations, we would therefore expect all subjects to choose the allocator group in the coordinator treatment despite their lack of knowledge, at the time of making the decision, of the subsequently resulting group size. We use this benchmark for the derivation of the second part of our first hypothesis:

Hypothesis 1b. *All subjects in the coordinator treatment select into the CCG.*

In the dictator treatment, we assert that a rational allocator would set the contribution of all group members equal to their endowments and set a contribution of zero for herself. This yields the following payoffs:

$$\pi_{i,DCG} = \begin{cases} E + \frac{\lambda}{n_{DCG}} \cdot E \cdot (n_{DCG} - 1) & \text{if subject } i \text{ is the dictator, and} \\ \frac{\lambda}{n_{DCG}} \cdot E \cdot (n_{DCG} - 1) & \text{if subject } i \text{ is not the dictator} \end{cases} \quad (9)$$

We reflect the incentives induced by this payoff function in the first part of our second hypothesis.

Hypothesis 2a. *In the DCG, dictators always contribute nothing themselves and the full endowment for all other group members.*

⁵ Note that *rational* in this case includes the assumption that the actors are only interested in maximizing their own payoff without regard for the payoff of others.

⁶ Strictly speaking, this is only true if subjects assign a positive probability to $n_{CCG} > 1$.

⁷ This result holds for any $\lambda > 1$ and, thus, for any public good.

Since every player who joins the DCG has a chance of $1/n_{DCG}$ to become the dictator, the conditionally expected payoff (assuming full contribution) of joining the DCG, given a group size of n_{DCG} , would be:

$$\mathbb{E}[\pi_{i,DCG}] = \frac{1}{n_{DCG}} \cdot E + \left(\frac{\lambda}{n_{DCG}} \cdot E \cdot (n_{DCG} - 1) \right) \quad (10)$$

where \mathbb{E} is the conditional expectations operator assuming equilibrium play (*i.e.*, full contribution in the ACG; no contribution in the VCG). It follows from Equations (2) and (4), as well as from our treatment of the special case of a group size of one that $\mathbb{E}[\pi_{i,DCG}] \geq \mathbb{E}[\pi_{i,VCG}]$, with $\mathbb{E}[\pi_{i,DCG}] = \mathbb{E}[\pi_{i,VCG}]$ iff $n_{DCG} = 1$. Thus, even though the resulting group sizes are as yet undetermined when subjects make their group choice, selecting into the VCG is nonetheless a dominant strategy. This is also the case for the worst possible outcome in the DCG when only two subjects join the DCG.⁸ This leads to the second part of our second hypothesis:

Hypothesis 2b. *All subjects in the dictator treatment select into the DCG.*

Despite their theoretical validity, we judge it likely that hypotheses 1a and 1b, as well as hypotheses 2a and 2b will not hold in our experiments. Experimental economists (among others) have shown that people do not behave in an exclusively payoff maximizing manner. In our setting, possible reasons for off-equilibrium behavior include the heterogeneity of social preferences, bounded rationality, salience effects, aversion to risk and/or losses and a dislike of competing *per se*. Unfortunately, there is a large number of different theories of, e.g., social preferences, such that it is not possible to include all of them with precise predictions. We will therefore formulate some hypotheses regarding expected deviations from perfectly rational behavior based on social value orientation and inequality aversion.

1.2. Social Preference Predictions

Social preference models (and also social value orientation) assume that individuals are not concerned about their own payoff alone, but also about the payoffs to others and the relative sizes of their own and others' payoffs. One specific form of social preferences is inequality-aversion. Outcome based models of inequality aversion assume that subjects are averse to differences in outcomes (see, e.g., [29,30]). This allows us to make a prediction regarding the differences in group choice between the coordinator and dictator treatments. Since no inequality is possible in the CCG, inequality aversion cannot be a cause for subjects choosing the VCG in the coordinator treatment. This is different in our dictator treatment. Here, the dictator can choose different contribution levels for himself and for the other DCG members, thereby increasing his payoff relative to the other group members'. This reduces the utility of inequality-averse subjects and renders the VCG relatively more

⁸ Following from Equation (4), this assertion implies the following inequality: $1/n_{ACG} \cdot E + \lambda/n_{ACG} \cdot E \cdot (n_{ACG} - 1) > E$. It is easy to show that it simplifies to $\lambda > 1$ if $E > 0$ and $n > 1$.

attractive to them.⁹ Since we assume that there likely are such subjects in our subject pool, we reflect this in our next hypothesis:

Hypothesis 3. *Subjects are more likely to choose the ACG in the coordinator treatment than in the dictator treatment.*

Once we drop the assumption of rational expectations, subjects can be assumed to update their expectations of other participants' behavior based on their observations of past outcomes. We expect an effect of the amount of the dictator contribution in the previous period on subjects' group choice in the subsequent period.

Hypothesis 4. *Subjects' likelihood of selecting into the DCG increases in the previous period's dictator contribution.*

As Equation (3) makes clear, the negative effects of low dictator contributions in the DCG are diluted with increasing group size, since the cost of dictator free riding is jointly borne by more group members. We expect that this dilution effect will make it more likely for dictator treatment subjects to join the DCG when they expect many others to do so.¹⁰ Note, however, that our subjects do not know the group size for the period for which they are currently making their group choice. We conjecture that they will use the group size information from the previous period as a proxy for the current period's DCG size when forming their expectations of the latter.¹¹

Hypothesis 5. *Subjects' likelihood of selecting into the DCG increases in the previous period's DCG size.*

We also explore the impact of allocators' social value orientation on their decisions in the experiment. Subjects with a pro-social value orientation not only care about themselves, but also about relevant others (see, e.g., [31,32]). On the other hand, pro-self individuals are more interested in their own payoff. Previous experiments show that a pro-social value orientation correlates with cooperative behavior in economic experiments (see, e.g., [33]).¹² On this basis, we expect pro-social dictators to set a higher contribution level for themselves than do pro-self dictators. In the CCG, pro-social and pro-self coordinators should behave the same way (setting the contribution of everyone equal to the endowment).

⁹ Note that dictators with social preferences (like inequality-aversion or deriving positive utility from the payoffs of others) might themselves contribute (fully) to the public good (see hypothesis 6). We thank an anonymous referee for pointing this out.

¹⁰ Note that the dilution effect is counteracted by the decrease in probability of being assigned the dictator role with the attendant higher possible payoff. Refer to the Appendix in Section 4 for a proof that the first effect outweighs the second. Strictly speaking, our argument is based on the net effect.

¹¹ While we do consider this question to be interesting, we did not judge it important enough to explicitly elicit group size expectations, which carries a risk of causing an experimenter demand effect.

¹² For a review, see [34].

Hypothesis 6. *Pro-social dictators set their own contribution higher than pro-self dictators.*

1.3. Hypotheses about Dynamics

In our final two hypotheses, we explore dynamic behavioral effects, which are not necessarily connected to social preferences. In particular, we expect the previous period's contribution behavior in the VCG to influence subjects' group choice. If there are high lagged contributions in the VCG, subjects may be induced to select into this group for two different reasons. First, with high contributions, this group becomes attractive for free riders who want to exploit the contributing members of the VCG. Second, pro-social subjects may be attracted by the high contributions, because they want to participate in the generation of a public good out of their own (and other group members') free decisions. This conjecture is founded in the work of Sen [35,36], who argues that the freedom of choice yields intrinsic value to humans. We summarize this line of reasoning in the following hypothesis:

Hypothesis 7. *Subjects' likelihood of selecting into the VCG increases in the previous period's average VCG contribution.*

Since high contributions in a group generally make this group more attractive, we finally also expect the contribution behavior in the ACG in the previous period to influence subjects' group choice. This leads to our final hypothesis:

Hypothesis 8. *Subjects' likelihood of selecting into the ACG increases in the previous period's average ACG contribution.*

2. Design and Procedures

Our experiment was part of a larger research program and comprised a total of 11 sessions with 12 subjects each. Sessions 2, 3, 4 and 6 used the coordinator treatment and sessions 7 through 13, the dictator treatment (sessions 1 and 5 used a design that is not the subject of this paper).¹³ The experiments were conducted at the Max-Jung laboratory at the University of Graz from April to July, 2012. The participants were recruited from a subject pool consisting mainly of students from the faculty of Social and Economic Sciences. The use of ORSEE [37] ensured that every subject could only participate in the experiment once. The experiment was programmed and conducted with z-Tree [38]. All payments were made in Euros, and the conversion rate from experimental currency

¹³ In sessions 7 through 10, a programming error caused the period end screen to display the current period's DCG contributions also for previous periods. It is for this reason that we conducted additional sessions. To compound this unfortunate streak, an irreparable server crash then forced us to terminate the experiment after the first round in session 11. The first round data from this session are unaffected, but no questionnaires or SVO measures were elicited. We perform robustness checks of all our results to control for possible effects from the programming error and server crash using the "clean" session 12 and 13 data as a benchmark. We find no material changes in our results. For this reason, we include the session 7 through 11 data in our analyses.

units to Euros was 25 ECU = 1 EUR. Average earnings were 11.90 EUR, including a show up fee of 2.5 EUR. On average, a session lasted 45 min.

2.1. Treatment Design

Subjects participate in two rounds of ten periods each of a public good game.¹⁴ Before making their contribution decisions, they chose one of two groups to join for the current period. In the coordinator (dictator) treatment, these are the VCG and the CCG (the VCG and the DCG). The choice of group impacts how subjects' contributions are determined. Furthermore, the earnings for a period depend only on the subjects in the same group. Once every subject has entered her contribution, the payoff for this period is calculated according to Equation (1). Note that all subjects are informed about the number of subjects in their group at the time of making their contribution decision. At the end of each period, subjects see a results screen, which they can study for a maximum of 60 s (other than for the results screen, there were no time limits anywhere in the experiment). There was no deception involved in the experiment.

2.1.1. Design Features Specific to the Coordinator Treatment

In the coordinator treatment, Equation (1) applies equally to subjects choosing the VCG and ones choosing the CCG. The difference to the VCG is that in the CCG, one subject is randomly chosen out of all group members to make the contribution decision for the entire group. This coordinator subject enters a contribution, which then applies to all CCG members, including the coordinator herself.

At the end of a period, subjects in the coordinator treatment see a results screen, which informs them about four parameters for each group (VCG and CCG), for the period just completed, as well as for all previous periods. These are (i) the number of subjects in the group, (ii) the average contribution in the group, (iii) the per capita earnings from the public good and (iv) the average ending wealth. In addition to this, they are informed about their personal starting endowment, their contribution, their return from the public good and their ending wealth.

2.1.2. Design Features Specific to the Dictator Treatment

In the dictator treatment, one subject is randomly chosen to assume the role of the dictator, similar to the case of the coordinator just described. However, in the DCG, the dictator subject enters two parameters. The first is the contribution that applies to all group members, except the dictator. The second is the contribution that applies only to the dictator. As indicated in Section 1, the dictator can, for example, choose to let all other subjects contribute their full endowment of 20 ECU while contributing nothing herself.

¹⁴ The ten periods of a round are treated as a logical unit and subjects are informed that "we will now move on to the second round consisting of ten periods", but the only difference between period 10 in round 1 and period 1 in round 2 is that the history of the first round's periods is no longer displayed on the results screen. The motivation for this design feature was to determine the extent of a possible restart effect.

At the end of a period, subjects in the dictator treatment see the results screen displayed in Figure 2. For the VCG, they learn the same parameters as subjects in the coordinator treatment. For the DCG, they are informed about (i) the number of subjects in the group, (ii) the contribution of the dictator, (iii) the contribution the dictator has chosen for all other DCG members, (iv) the per capita earnings from the public good and (v) the average ending wealth. Furthermore, they also receive information about their own starting endowment, contribution, return from the public good and ending wealth.

2.2. Session Structure

At the beginning of an experimental session, subjects arrive and wait outside the laboratory. At the designated starting time, subjects are welcomed by the experimenter, draw cards with their computer numbers, are led into the lab and sit down at the workstations corresponding to the numbers on their cards. They there find a printed set of instructions, which the experimenter reads out loud, asking the subjects to read along. After answering any possible remaining questions individually, the experimenter then starts the first round of 10 experimental periods.

Figure 2. Example dictator treatment results screen. The figure displays an example of the results screen as shown to subjects at the end of each period in the dictator treatment. Text printed in red (grey in greyscale printouts) is a translation of the original German captions.

Ergebnis Periode 2 Result Period 2										
Durchschnittliche Gruppenergebnisse Average Group Results										
Gruppe mit individueller Beitragswahl Group with Individual Contribution Choice					Gruppe mit Beitragswahl durch zufällig bestimmten Spieler Group with Contribution Choice by Randomly Chosen Player					
Periode	Spielerzahl	Durchschnittl. Beitrag/Kopf	Projektertrag/Kopf	Durchschnittl. Endpunktestand	Periode	Spielerzahl	Beitrag zufällig best. Spieler	Beitrag andere Spieler	Projektertrag/Kopf	Durchschnittl. Endpunktestand
1	3	5.33	8.53	23.20	1	9	5	19	27.91	28.91
2	2	3.00	4.80	21.80	2	10	0	20	28.80	28.80
Period	Number of Players	Average per Capita Contribution	Project Return per Capita	Average Ending Account Total	Period	Number of Players	Contribution Randomly Chosen Player	Contribution Other players	Project Return per Capita	Average Ending Account Total
Ihr Ergebnis Your Result										
You were a member of the group with... Sie waren Mitglied der Gruppe mit...					Contribution Choice by Randomly Chosen Player Beitragswahl durch zufällig bestimmten Spieler					
Your Starting Account Total: Ihr Startpunktestand:					20.00 Punkte 20.00 Points					
Your Contribution: Ihr Beitrag:					20 Punkte 20 Points					
Your Earnings: Ihr Ertrag:					28.80 Punkte 28.80 Points					
Your Ending Account Total: Ihr Endpunktestand:					28.80 Punkte 28.80 Points					

In each period, subjects can first choose the group they want to join. If they join the ACG, they then learn whether they have been randomly chosen for the role of the allocator in this period. Following this, coordinator subjects enter the contribution they want every CCG member (including themselves) to make. Dictator subjects enter both a contribution they want to make themselves and a

contribution they want their fellow DCG members to make. Subjects in the VCG enter the contribution they want to make. Once all contribution decisions have been entered, a results screen informs subjects of the outcomes of the present and previous periods in the current round. After the first round is over, the experimenter starts the second round, which proceeds analogously to the first. After the second round, one of the two is randomly selected for payoff in order to avoid portfolio effects. This is achieved by letting one of the participants publicly throw a die, the result of which determines the payoff-relevant round for all subjects. Once the experimenter has entered this information into his computer, subjects are informed about their payoffs on their screens.

The experimenter then starts a computerized questionnaire eliciting data on subject demographics and on their experiences and strategies in the experiment. The experimenter furthermore hands out a sheet for the elicitation of the SVO, which he asks every subject to fill in. Subjects who have finished filling in the questionnaire and SVO sheet [39] step outside the lab to wait until everybody else has also finished. Once this is the case, the experimenter asks subjects to step into the lab one at a time and pays them anonymously.

3. Results

3.1. Summary Statistics

We begin by discussing the summary statistics listed in Table 1. The first block of three rows already lets us reach a verdict regarding our hypotheses 1b and 2b. A one sample Wilcoxon signed-rank test confirms that the median CCG (DCG) size is significantly different from 12 (both p -values: 0.000). We can repeat this type of analysis for the second block of three rows in Table 1 to obtain results for hypotheses 1a and 2a. The median contribution in the CCG is significantly different from 20 (one-sample Wilcoxon signed-rank test p -value: 0.000). Similarly, we can reject both conjectures in hypothesis 2a: dictators' median contributions for themselves and for the other DCG members are significantly different from zero and 20, respectively (both one-sample Wilcoxon signed-rank test p -values: 0.000).

Next, we look at key behavioral patterns in the CCG and the DCG. In the CCG, we observe full contribution by all group members in 85% of all cases. This is a value higher than usually observed in public good games. On the other hand, this means that in 15% of all periods, the coordinator did not choose the full contribution for all members.¹⁵ This is surprising, since full contribution by coordinators maximizes both individual and collective welfare. Possible explanations for this seemingly irrational behavior include antisocial preferences, subject confusion or boredom and the wish to generate an outcome differing from the usual pattern. While we find no correlation between coordinators not fully contributing and their SVO values, thus speaking against the first explanation, further experiments would be needed to determine what is responsible for this inefficiency.

¹⁵ Note that the CCG (DCG) had a minimum of six (four) members in all periods.


Table 1. Descriptive statistics. The table shows summary statistics for the coordinator (C) and dictator (D) treatments, separately for the voluntary contribution groups (VCG) and the allocator contribution groups (coordinator contribution group (CCG); and dictator contribution group (DCG)). Note that in the case where we report on the contributions for the first and the last period in the ACGs, the sample size is rather small, with seven observations for the DCG and four observations for the CCG.

Group	VCG (C)		VCG (D)		CCG		DCG	
	1	2	1	2	1	2	1	2
Average group size n	2.6	1.98	2.91	1.8	9.4	10.03	9.09	10.20
Average contribution	8.59	6.84	6.97	6.65	18.00	18.98	17.39	18.50
Average dictator contribution							10.40	10.22
Average group member contribution							18.26	19.40
Average contribution period 1	7.88	5.23	6.90	5.36	20.00	13.71	14.93	16.51
Average contribution period 10	3.43	10.00	5.62	4.75	19.41	20.00	17.86	18.08
% zero contribution period 1	6.30	38.5	5.00	27.30	0.00	31.40	2.30	0.00
% zero contribution period 10	57.10	0.00	38.50	50.00	0.00	0.00	2.80	3.30
% full contribution period 1	12.50	7.70	10.00	9.10	100.00	68.60	63.60	50.80
% full contribution period 10	14.30	0.00	7.70	8.30	80.50	100.00	83.10	75.00

In the DCG, the share of 41% full contributions by all group members is lower than in the CCG. This is of course an effect of our treatment differences, which allowed our dictators to let all other DCG members contribute fully, while themselves free riding. This happened in 26% of the periods. In the remaining periods, we observe other behavioral patterns. In 16% of the periods, the dictator chose full contribution for all other DCG members and contributed an amount larger than zero, but smaller than the endowment herself. In 17% of the periods, the dictator chose a contribution lower than the endowment for the other DCG members.

Figure 3 shows the mean contributions for both groups and both treatments. The observed behavioral patterns result in significantly higher contributions in both the DCG and the CCG than in the VCG (both two-sided Wilcoxon rank-sum test p -values: 0.000). As shown in Table 1, we observe a median contribution of 20 for the first and the last periods of both rounds in both the CCG and the DCG. On the other hand, we observe median contributions between zero and ten in the VCG in both treatments.


Figure 3. Contributions. Mean period contributions by treatment and group. The data of rounds 1 and 2 are pooled.



3.2. Group Choice

We continue our analysis with the average group sizes and their development over time. Figure 4 shows that the average group size of the allocator group starts at a very high level and, over time, increases even further in both treatments. Consequently, only a few subjects join the VCG.

Figure 4. Overall trend in average group size over time in the two treatments. The figure displays the average group size in each of the ten periods separately for the dictator and the coordinator treatments. The data of rounds 1 and 2 are pooled. The dashed lines are linear predictions. The thin lines are average group sizes in individual sessions.



Next, we conduct a regression analysis of the aggregate data, which is presented in Table 2. In the models (which we fit individually for each treatment), we control for the round, the period and the interaction between the two, allowing the slope of the periods to vary between rounds. We also include previous period data on the average size of, and the average contribution in, the ACG.

Table 2. Group size. The table shows the results of OLS regressions of the CCG and DCG group sizes on a number of regressors. Round is 1 (2) in the first (second) of the two 10-period sequences. Period2 equals the period number Period in round 2 and zero otherwise. The remaining variables are the lagged average contributions in the ACG and the lagged ACG group size (of the DCG in the regression of the dictator treatment, of the CCG in the analysis of the coordinator treatment) and the lagged contribution of the dictator in the DCG.

Regressors	Model 1	Model 2
	CCG Group Size	DCG Group Size
Round	0.098 (0.538)	0.690 (0.516)
Period	0.069 (0.071)	0.060 (0.056)
Period2	0.041 (0.104)	-0.078 (0.069)
Lagged Average Contribution in ACG	0.103 (0.016) ***	0.089 (0.034) **
Lagged Dictator Contribution in DCG		0.033 (0.014) *
Lagged Group Size in ACG	0.287 (0.157)	0.393 (0.059) ***
Constant	4.537 (0.937) **	2.942 (0.889) **
R^2	0.30	0.37
Adjusted R^2	0.25	0.34
N	72	117

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Standard errors clustered at the session level (in parentheses).

Our results show that the time effect visible in Figure 4 is not significant in either treatment when controlling for the variables that were identified as relevant in our hypotheses in Section 1. Models 1 and 2 share one statistically significant effect: a larger average ACG contribution in the previous round results in a larger current ACG group size, supporting our hypothesis 8. It also implies that when the allocator contributes relatively little, the group size of the VCG increases in the following period. Furthermore, for the dictator treatment only, we find a significant effect of the DCG group size in the previous period. This provides support for the presence of the dilution effect, as conjectured in hypothesis 5. Finally, we find a significant effect of the dictator contribution in the previous period on the DCG group size (hypothesis 4).¹⁶

¹⁶ Tobit regression censored at zero and 12 yields similar results.

We continue our analysis by investigating individual subjects' group choice behavior using Probit models.¹⁷ Table 3 presents the two regression models we believe best reflect the structural relationships in our data. In Model 3, we include the Round and Period variables, a treatment dummy and a variable containing our subjects' social value orientation as measured using the instrument defined in Murphy *et al.* [28]. Higher values of the SVO measure indicate a greater willingness to give up own income to benefit others. We also include an interaction of SVO with the treatment dummy, as well as variables containing information on the previous period's average contributions in the VCG and ACG (AvgContribVCG_L and AvgContribACG_L) and on the size of the ACG (and, separately, for the DCG) in the previous period (GroupSizeACG_L and GroupSizeACG_L_x_Dictator).

As predicted in hypothesis 3, we find a negative treatment effect. Subjects are less likely to join the ACG in the dictator treatment. Our model also shows that the likelihood of a subject selecting into the ACG as opposed to the VCG increases in the average contribution in the ACG in the previous period, thus lending further support to our hypothesis 8. Conversely, higher contributions in the VCG in the previous period decrease the likelihood of subjects joining the ACG. This result supports hypothesis 7. We also find that greater social value orientation of a subject cannot be shown to play a role.¹⁸ Finally, the results do not yield any evidence of subjects being more likely to select into the ACG in later periods or in the second round.

Model 4 differs from Model 3 in that it uses only the dictator treatment data and contains the amount of the previous period's dictator contribution as an additional variable. We find a statistically significant effect of the lagged size of the DCG, but not the ACG in general, lending further support to the dilution effect of hypothesis 5. Model 4 also shows a significant effect of the lagged average contribution in the DCG, but none of the lagged dictator contribution. Hypothesis 4 does not receive support from this result. The average contribution in the VCG in the previous period shows a highly significant negative coefficient, again, in line with our hypothesis 7. Just as in the pooled analysis, there appears not to be a material effect of SVO on group choice.

¹⁷ Robustness checks confirm that our main results are stable with regard to the use of a logit model and to the inclusion or exclusion of different questionnaire items. A translation of the questions is provided in the Electronic Supplementary Information.

¹⁸ Two subjects did not fill in the SVO questionnaire correctly and are, therefore, excluded from all analyses employing social value orientation data. Furthermore, the session 11 data are not included, since no SVO or other questionnaire items were elicited, due to the server crash.

Table 3. Determinants of allocator group choice. The table shows the results of Probit regressions of IsACG on a number of regressors. IsACG is a dummy equal to zero (one) if the subject chooses the VCG (ACG) in a period. Dictator is a dummy variable equal to zero (one) in the coordinator (dictator) treatment. Round is 1 (2) in the first (second) of the two 10- period sequences. Period2 equals the period number Period in round 2 and zero otherwise. SVO is the social value orientation of the subjects, measured using the slider-measure from Murphy *et al.* (2011), and SVO x Dictator is SVO interacted with Dictator. The remaining variables are the lagged average contributions in the ACG and VCG, the lagged ACG group size (pooled and in the dictator treatment only) and the lagged contribution of the dictator in the dictator treatment (note that the model that includes this variable solely uses dictator treatment data). Logit estimation and logit panel regression yield similar results. Robustness checks, where we include various questionnaire response items and interaction terms, yield no clear effects from the control variables, but leave the main effects unchanged.

Regressors	Model 3	Model 4
Subsample	All cases	Dictator treatment
Dictator	-1.284 (0.679) *	
Round	-0.127 (0.187)	-0.012 (0.303)
Period	0.000 (0.017)	0.008 (0.026)
Period2	0.026 (0.027)	0.010 (0.041)
SVO	-0.010 (0.011)	
SVO_x_Dictator	0.020 (0.015)	0.010 (0.010)
AvgContribACG_L	0.030 (0.008) ***	0.024 (0.011) **
AvgContribVCG_L	-0.046 (0.012) ***	-0.038 (0.012) ***
GroupSizeACG_L	0.044 (0.061)	
GroupSizeACG_L_x_Dictator	0.085 (0.066)	0.119 (0.038) ***
DictContrib_L		0.009 (0.007)
Constant	0.661 (0.594)	-0.750 (0.475)
Pseudo R^2	0.04	0.04
N	1,358	746

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Standard errors clustered at the subject level (in parentheses).

3.3. Contributions

We next focus on dictators' contribution behavior in the DCG. Table 4 contains our OLS regression results for the dictator's own contribution (Model 5). We control for a time trend, the lagged contributions, the group size of the ACG, the lagged dictator contribution and SVO. First, this is the only instance where we detect a significant influence of SVO—in this case, the SVO of the dictator subject—on experimental behavior. This lends support to our hypothesis 6. Second, the dictator's own contribution increases in the previous period's dictator contribution.

Table 4. Determinants of dictator contribution choice. The table shows the results of an OLS regression of the dictator’s own contribution on a number of regressors. Round is 1 (2) in the first (second) of the two 10-period sequences. Period2 equals the period number Period in round 2 and zero otherwise. SVO is the social value orientation of the dictator, measured using the slider-measure from Murphy *et al.* [28]. The remaining variables are the lagged average contribution in the DCG, the concurrent DCG group size and the lagged dictator contribution. The inclusion of the lagged variables leads to the exclusion of the observations from period 1 in rounds 1 and 2. Tobit regression censored at zero and 20 yields similar results. Robustness checks, where we include various questionnaire response items and interaction terms, yield no clear effects from the control variables, but leave the main effects unchanged.

Regressors	Model 5
Round	-2.225 (4.214)
Period	-0.691 (0.454)
Period2	0.417 (0.632)
AvgContribDCG_L	-0.159 (0.266)
GroupSizeDCG	0.736 (0.486)
DictatorContribution_L	0.325 (0.104) ***
SVO	0.127 (0.060) **
Constant	5.132 (7.438)
R^2	0.23
Adj. R^2	0.17
N	107

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Standard errors clustered at the subject level (in parentheses).

4. Discussion and Conclusion

The experimental literature on mechanisms fostering cooperation in dilemmas has lately predominantly focused on the effectiveness of punishment and reward. Recent work by Hamman *et al.* [20] and Bolle and Vogel [21] has extended this field to encompass inequalities in the power over one’s own and others’ decisions. Such asymmetries in the decision-making powers of economic actors are a frequent and important phenomenon outside the laboratory and, as such, merit careful analysis.

Hamman *et al.* [20] and Bolle and Vogel [21] demonstrate the possible efficiency gains from centralized decision-making in the provision of a public good. We extend their research by analyzing allocator mechanisms with different action spaces. We also implement direct competition between different contribution mechanisms by allowing for endogenous group choice and investigate to what extent social preferences drive contribution and group choice behavior. We find that the vast majority of our subjects is willing to cede decision authority to a central planner in

order to reap efficiency gains from improved coordination.¹⁹ We consider this result of great importance, since it clearly shows that human subjects are willing to submit to a randomly selected centralized authority if it leads to higher (expected) average payoffs. This is true both in a setting enforcing equality in payoffs and in one where the subject endowed with decision authority for the entire group can exploit this power to maximize her own payoffs at the expense of her team members'. Nonetheless, subjects are more likely to select into the allocator group in the first than in the second setting, and we investigate the factors driving this decision. Our data shows that subjects condition their group choice on historical group sizes and contribution behavior. Finally, we find that an allocator's social value orientation plays a role in her contribution choice in the setting where she has the option to exploit her fellow group members.

We summarize our findings in Table 5. Overall, they show that the allocator mechanism is not only more successful in establishing high contributions than a voluntary contribution scheme, but also wins out in a direct competition. We believe that these encouraging results merit further research into allocator contribution mechanisms for the provision of public goods and the accompanying power asymmetries.

Table 5. Overview of hypotheses and results. The table shows all hypotheses derived in Section 1 and the corresponding results from Section 3.

No.	Statement	Result
1a	In the CCG, coordinators always contribute the full endowment.	Rejected; only in 85% of cases
1b	All subjects in the coordinator treatment select into the CCG.	Rejected; median group size significantly smaller than 12
2a	In the DCG, dictators always contribute nothing themselves and the full endowment for others.	Rejected; only in 26% of cases
2b	All subjects in the dictator treatment select into the DCG.	Rejected; median group size significantly smaller than 12
3	Subjects are more likely to choose the CCG than they are to choose the DCG.	Supported; see Model 3 in Table 3
4	Subjects' likelihood of selecting into the DCG increases in the previous period's dictator contribution.	Not supported; see Model 4 in Table 3
5	Subjects' likelihood of selecting into the DCG increases in the previous period's DCG size.	Supported; see Models 2 and 4 in Tables 2 and 3, respectively
6	Pro-social dictators set their own contribution higher than pro-self dictators.	Supported; see Model 5 in Table 4
7	Subjects' likelihood of selecting into the VCG increases in the previous period's average VCG contribution.	Supported; see Models 3 and 4 in Table 3

¹⁹ Note that the allocator mechanism at the same time decreases the risk of exploitation and impacts payoff inequality. Further research is needed to disentangle the differential effect of these factors.

Acknowledgments

We thank two anonymous referees and the audiences at the International Meeting of the Economic Science Association 2012, the Annual Meeting of the Austrian Economic Society 2013, the 15th International Conference on Social Dilemmas, the annual conference of the GfeW 2012, the Quantitative Sociology Colloquium at the ETH Zürich and the research seminars at the Vienna Center of Experimental Economics and at the Faculty of Social and Economic Sciences at the University of Graz for valuable comments. All errors remain our own.

Conflicts of Interest

The authors declare no conflict of interest.

Appendix

Proof of the Dilution Effect in the Dictator Treatment

Remember that Equation (4) posited the following expected payoff in the DCG:

$$\mathbb{E}[\pi_{i,DCG}] = \frac{1}{n_{DCG}} \cdot E + \left(\frac{\lambda}{n_{DCG}} \cdot E \cdot (n_{DCG} - 1) \right)$$

Proof. For the dilution effect to obtain, the following inequality must then hold (for expositional convenience, we suppress the DCG subscript):

$$\frac{1}{n} \cdot E + \left(\frac{\lambda}{n} \cdot E \cdot (n - 1) \right) < \frac{1}{n + c} \cdot E + \left(\frac{\lambda}{n + c} \cdot E \cdot (n + c - 1) \right) \quad (5)$$

where c is an arbitrary, positive integer. Inequality (5) can be simplified to:

$$c < \lambda c \quad (6)$$

Inequality (6) is fulfilled for any $n \in \mathbb{N}^+$, $\lambda > 1$ and $E > 0$ and, thus, for any public good. This also extends to our parameterization, where $\lambda = 1.6$. □

References

1. Kennedy, D.; Norman, C. What don't we know? *Science* **2005**, *309*, 75–75.
2. Pennisi, E. How did cooperative behavior evolve? *Science* **2005**, *309*, 93–93.
3. Ledyard, J. Public Goods: A Survey of Experimental Research. In *Handbook of Experimental Economics*; Kagel, J., Roth, A., Eds.; Princeton University Press: Princeton, NJ, USA, 1995.
4. Chaudhuri, A. Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature. *Exp. Econ.* **2011**, *14*, 47–83.
5. Keser, C.; van Winden, F. Conditional cooperation and voluntary contributions to public goods. *Scand. J. Econ.* **2000**, *102*, 23–39.

6. Fischbacher, U.; Gächter, S.; Fehr, E. Are people conditionally cooperative? Evidence from a public goods experiment. *Econ. Lett.* **2001**, *71*, 397–404.
7. Fehr, E.; Fischbacher, U. The nature of human altruism. *Nature* **2003**, *425*, 785–791.
8. Fehr, E.; Gintis, H. Human motivation and social cooperation: Experimental and analytical foundations. *Ann. Rev. Sociol.* **2007**, *33*, 43–64.
9. Ostrom, E.; Gardner, R.; Walker, J. *Rules, games, and common-pool resources*; University of Michigan Press: Ann Arbor, MI, USA, 1994.
10. Fehr, E.; Gächter, S. Fairness and retaliation: The economics of reciprocity. *J. Econ. Perspect.* **2000**, *14*, 159–181.
11. Fehr, E.; Gächter, S. Altruistic punishment in humans. *Nature* **2002**, *415*, 137–140.
12. Walker, J.M.; Halloran, M.A. Rewards and sanctions and the provision of public goods in one-shot settings. *Exp. Econ.* **2004**, *7*, 235–247.
13. Sefton, M.; Shupp, R.; Walker, J.M. The effect of rewards and sanctions in provision of public goods. *Econ. Inq.* **2007**, *45*, 671–690.
14. Margreiter, M.; Sutter, M.; Dittrich, D. Individual and collective choice and voting in common pool resource problem with heterogeneous actors. *Environ. Resour. Econ.* **2005**, *32*, 241–271.
15. Walker, J.M.; Gardner, R.; Herr, A.; Ostrom, E. Collective choice in the commons: Experimental results on proposed allocation rules and votes. *Econ. J.* **2000**, *110*, 212–234.
16. Weber, M.; Kalberg, S. *Max Weber: Readings and commentary on modernity*; Blackwell Publisher: Malden, MA, USA, 2005.
17. Mann, M. *A History of Power from the Beginning to A.D. 1760*, reprinted ed.; Cambridge University Press: Cambridge, UK, 1997.
18. Fehr, E.; Herz, H.; Wilkening, T. The lure of authority: Motivation and incentive effects of power. *Am. Econ. Rev.* **2013**, *103*, 1325–1359.
19. Kroll, S.; Cherry, T.L.; Shogren, J.F. Voting, punishment, and public goods. *Econ. Inq.* **2007**, *45*, 557–570.
20. Hamman, J.R.; Weber, R.A.; Woon, J. An experimental investigation of electoral delegation and the provision of public goods. *Am. J. Polit. Sci.* **2011**, *55*, 738–752.
21. Bolle, F.; Vogel, C. Power comes with responsibility—or does it? *Public Choice* **2011**, *148*, 459–470.
22. Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action*; Cambridge University Press: Cambridge, UK and New York, NY, USA, 1990.
23. Sutter, M.; Haigner, S.; Kocher, M.G. Choosing the carrot or the stick? Endogenous institutional choice in social dilemma situations. *Rev. Econ. Stud.* **2010**, *77*, 1540–1566.
24. Gürrer, Ö.; Irlenbusch, B.; Rockenbach, B. The competitive advantage of sanctioning institutions. *Science* **2006**, *312*, 108–111.
25. Kosfeld, M.; Okada, A.; Riedl, A. Institution formation in public goods games. *Am. Econ. Rev.* **2009**, *99*, 1335–1355.
26. Service, E.R. *Primitive Social Organization: An Evolutionary Perspective*; Random House: New York, NY, USA, 1962.

27. Rockenbach, B.; Milinski, M. The efficient interaction of indirect reciprocity and costly punishment. *Nature* **2006**, *444*, 718–723.
28. Murphy, R.O.; Ackermann, K.A.; Handgraaf, M. Measuring social value orientation. *Judgm. Decis. Mak.* **2011**, *6*, 771–781.
29. Bolton, G.; Ockenfels, A. ERC: A theory of equity, reciprocity, and competition. *Am. Econ. Rev.* **2000**, *90*, 166–193.
30. Fehr, E.; Schmidt, K. A theory of fairness, competition, and cooperation. *Q. J. Econ.* **1999**, *114*, 817–868.
31. Murphy, R.O.; Ackermann, K.A. Social value orientation: Theoretical and measurement issues in the study of social preferences. *Personal. Soc. Psychol. Rev.* 2013, doi:10.1177/1088868313501745.
32. Ackermann, K.A.; Fleiß, J.; Murphy, R.O. Reciprocity as an individual difference; Working Paper Series, Social and Economic Sciences 2013-05, Faculty of Social and Economic Sciences, Karl-Franzens-University Graz, Austria, 2013.
33. Fleiß, J.; Leopold-Wildburger, U. Social Values and Cooperation. Results from an Iterated Prisoner's Dilemma Experiment. In *Operations Research Proceedings 2011*; Klatte, D., Lüthi, H.J., Schmedders, K., Eds.; Springer Verlag: Berlin/Heidelberg, Germany, 2012.
34. Balliet, D.; Parks, C.; Joireman, J. Social value orientation and cooperation in social dilemmas: A meta-analysis. *Group Process. Intergroup Relat.* **2009**, *12*, 533–547.
35. Sen, A. Welfare, preference and freedom. *J. Econo.* **1991**, *50*, 15–29.
36. Sen, A. Freedom of choice: Concept and content. *Eur. Econ. Rev.* **1988**, *32*, 269–294.
37. Greiner, B. An Online Recruitment System for Economic Experiments. In *Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht 63*; Kremer, K., Macho, V., Eds.; Ges. für Wiss. Datenverarbeitung: Göttingen, Germany, 2004.
38. Fischbacher, U. z-Tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* **2007**, *10*, 171–178.
39. Chair of Decision Theory and Behavioral Game Theory. Paper based SVO Slider measure. Available online: http://vlab.ethz.ch/svo/SVO_Slider/SVO_Slider.html (accessed on 10 October 2013).

Introducing Disappointment Dynamics and Comparing Behaviors in Evolutionary Games: Some Simulation Results

Tassos Patokos

Abstract: The paper presents an evolutionary model, based on the assumption that agents may revise their current strategies if they previously failed to attain the maximum level of potential payoffs. We offer three versions of this reflexive mechanism, each one of which describes a distinct type: spontaneous agents, rigid players, and ‘satisficers’. We use simulations to examine the performance of these types. Agents who change their strategies relatively easily tend to perform better in coordination games, but antagonistic games generally lead to more favorable outcomes if the individuals only change their strategies when disappointment from previous rounds surpasses some predefined threshold.

Reprinted from Special Issue: Aspects of Game Theory and Institutional Economics, *Games*. Cite as: Patokos, T. Introducing Disappointment Dynamics and Comparing Behaviors in Evolutionary Games: Some Simulation Results. *Games* **2014**, *5*, 1–25.

1. Introduction

Individuals are averse to unpleasant experiences. When such experiences happen, it makes sense to assert that the individuals affected shall choose different strategies from those that brought about unsatisfactory outcomes. In this paper we present three variations of a learning model, which reflects the intuitive fact that agents try to eschew perceived disappointing outcomes. Disappointment emerges when a player fails to achieve the maximum level of payoffs that would be possible, had a different strategy been chosen. In other words, the core assumption is that a differential between someone’s actual payoffs and the maximum level of potential payoffs (with the opponent’s choice taken as a given) generates a tendency to choose a different strategy in the next round of the same game.

While it seems safe to conjecture that individuals avoid disappointment in general, individual reactions to past disappointment outcomes are contingent on psychological issues, and, as such, they may vary dramatically across persons. For example, a person with relatively low tolerance to disappointment might be expected to change their strategy after a disappointing outcome with higher probability than another person who is more patient. Evidently, the individual psychological profile is important in determining action. Each one of the three variations of the learning model we describe represents a distinct behavioral type: we deal with spontaneous and impatient agents, players who are rigid and display high inertia, and ‘satisficers’. We use a simulation program to study the evolutionary equilibria in an assortment of 2×2 games with populations consisting of the aforementioned behavioral types, the aim being to compare these types’ performances in coordination and antagonistic games.

Several prominent authors (such as Sugden [1] or Rubinstein [2]) have expressed the opinion that economics of bounded rationality does not need more theoretical models, but rather, must focus on

empirical research. While we embrace this view, the underlying behavioral hypotheses we present here do not seek to explain specific empirical or experimental findings; however, our purpose is not to merely enrich the literature with more learning or adaptation rules, which would perhaps seem superfluous. Rather, we focus on providing a better handle on how a combination of limited computational power and of a psychological aversion to disappointment matters. Although the rules we use are *ad hoc* constructed, they serve as proxies of players' real-life attitudes (for example, patience or spontaneity). In fact, our simulation results show that, depending on the game, different behavioral types are bound to be more successful than others; this means that a single reflexive or adaptation model is most likely to be insufficient for studying interactions, which reinforces the need for more empirical research.

One might wonder why we choose to introduce newer learning rules, rather than use something from the wealth of rules that can be found in the literature. The three rules presented in this paper have a clear behavioral background, based on the postulation that what determines human action now is possible disappointment experienced in past play. Therefore, even if the adaptive procedures we suggest could be approximated by an existing rule, we are interested in exploring how the *specific* stylized assumptions (each one of which is linked to a psychological type) translate to a learning model, and at a second level, in comparing these particular rules as to their social efficiency; it is then possible to conclude if these characteristics help the individuals attain better outcomes (such as a peaceful resolution in the “Hawk–Dove” game, or mutual cooperation in the “Prisoners’ Dilemma”). The corresponding findings, even if they are on a theoretical level, are important to those interested in the psychology of strategic behavior, in that they provide newer insights on intertemporal strategic play; and if the behavioral type can be seen as a control variable (*i.e.*, if the individual may choose their own type), then this theoretical approach is apt to suggest what type would be preferable, given the nature of the interaction.

The paper is structured in five sections. Section 2 provides a concise review of the use of adaptive procedures in evolutionary game theory, along with the notion of disappointment in economics. Section 3 presents the three different agent type protocols. Section 4 discusses simulations with and without random perturbations (*i.e.*, noise on the Markovian process), and Section 5 concludes.

2. Evolutionary Dynamics and the Notion of Disappointment

Evolutionary game theory took off in the first half of the 1970s, after Maynard Smith and Price [3] and Maynard Smith [4], inspired by Lewontin [5], applied game theory to biology and defined the evolutionary stable strategy. Taylor and Jonker [6] proposed replicator dynamics as a way to translate this process to mathematical language. Although very popular (mainly due to their simplicity), replicator dynamics are not generally thought of as particularly apt for economic applications¹ because they are seen as a highly restrictive selection model. Interest, therefore, shifted to dynamics that seek to describe how a population increases along with the success of its adopted strategy. According to Friedman [11], the most abstract way of modeling a selection process is to assume growth rates that are positively correlated with relative fitness.

¹ For comprehensive reviews, see [7–8]. More recent texts include [9–10].

Usually, evolutionary selection processes do not just portray how a population grows over time, but they are also based on an underlying series of assumptions that offer an explicit rationale of how agents adapt with historical time. Except for natural selection, Young (1998) distinguishes between rules of imitation (for example, [12]), reinforcement learning [13–15], and fictitious play (for example, [16]). In some of these models, beliefs on what the opponent plays are updated by means of Bayes' rule (see [17]).

Evolutionary dynamics can be deterministic or stochastic. Deterministic evolutionary dynamics (such as the replicator dynamics) describe the evolutionary path with systems of differential (or difference) equations; each one of these equations expresses the increasing (or decreasing) rate of some strategy's frequency as a function of the portion of the population who chooses this strategy, and in accordance to the assumed revision protocol [18–19]. In such models, mutations are thought of as rare and random, and therefore, not continual or correlated in any way. Contrary to this assumption, stochastic evolutionary dynamics examine the possibility that mutations cause enough "noise" to pull the population state out of a certain basin of attraction. These models incorporate perturbations by use of Markov processes and may lead to fundamentally different results than deterministic dynamics. Among the seminal works in stochastic dynamics are Foster and Young [20], Kandori, Mailath, and Rob [21], Young [22], and Binmore, Samuelson, and Vaughan [23]. A comprehensive presentation is offered in Young [24].

At first, this paper describes an agent-driven, stochastic evolutionary model featuring heuristic learning. In each new period, players observe whether their past choices have been best replies to those of their opponents. If not, then "disappointment" emerges, and the players present a tendency to switch to alternative strategies in the future. Individual choice based on the avoidance of disappointment or regret has been discussed in the literature by various authors [25–29]. A comparative presentation of disappointment and regret aversion models appears in Grant *et al.* [30]. Our paper presents a few stylized, novel ways to translate the core idea into an evolutionary game theoretic context.

In Loomes and Sugden seminal contribution [25], agents experience regret when some choice they made proves to be less successful (in payoff terms) than another option they did not choose; individuals are aware of this effect, and make their choices so as to avoid regret. In our model, individuals, being boundedly rational, do not take proactive measures to circumvent regret, but rather, they act upon realized disappointing outcomes. The revision protocols that we use to implement this idea are close in character to the adaptive procedure studied in Hart and Mas-Colell [31], where the concept of '*regret matching*' is introduced. While regret matching requires quite demanding computational abilities from the part of the agents, the model presented here is more heuristic and deals with less sophisticated players (not necessarily in the sense that they are less smart, but mainly because they have shorter memory and lack perfect knowledge of their surroundings).

The paper explores these dynamics by use of simulation software, confining the analysis to 2×2 symmetric games. The aim is to gain insights on the properties of the different revision protocols (each one of which corresponds to a specific behavioral type), and, thus, investigate the possibility that some of these types consistently outperform the rest. The software also allows for stochastic shocks in the form of 'intruders' (*i.e.*, preprogrammed automata), who may be matched, with positive

probability, with members of the original population. Allowing for such perturbations is important, because they account for random errors or even calculated deviations (or interventions), which can potentially affect the evolutionary course, often in unexpected ways. Simulation software programs are, in any case, being used increasingly in the literature in the study of evolutionary games; see, for instance, [32–35].

3. The Model

3.1. Theoretical Background

Let $I = \{1, 2, \dots, N\}$ be the set of players, and for each player i , let S_i be her finite set of pure strategies. Following Weibull [36], each player's pure strategies are labeled by positive integers. Thus, $S_i = \{1, 2, \dots, m_i\}$, for some integer $m_i > 1$. The set of pure strategies in the game, denoted S , is the Cartesian product of the players' pure strategy sets. For any strategy profile $s \in S$ and player $i \in I$, let $\pi_i(s) \in \mathbb{R}$ be the associated payoff to player i . Let $\pi: S \rightarrow \mathbb{R}^n$ be the combined pure-strategy payoff function of the game (that is, the function assigning to each pure-strategy profile the vector of payoffs $(\pi_1(s), \pi_2(s), \dots, \pi_n(s))$). With the above notation, the game is summarized by the triplet $G = (I, S, \pi)$.

We use s^{-i} to denote a pure strategy combination of all players except i and define the set of i 's pure best replies to a strategy combination s^{-i} to be the nonempty finite set $BR_i = \{h \in S_i : \pi_i(h, s^{-i}) \geq \pi_i(k, s^{-i}) \forall k \in S_i\}$. We define the set of i 's pure worst replies to a strategy combination s^{-i} to be the nonempty finite set $WR_i = \{w \in S_i : \pi_i(w, s^{-i}) \leq \pi_i(k, s^{-i}) \forall k \in S_i\}$. We define the *maximum disappointment* $\Delta_i \in \mathbb{R}$ as $\Delta_i = \max\{\pi_i(h, s^{-i}) - \pi_i(w, s^{-i})\}$, for all s^{-i} , $h \in BR_i$, $w \in WR_i$. In other words, for each possible s^{-i} , we calculate the difference between the payoffs associated with the best reply and the payoffs associated with the worst reply to s^{-i} ; thus, we have $m_1 m_2 \dots m_{i-1} m_{i+1} \dots m_n$ non-negative real numbers. Δ is the maximum of these numbers.

Suppose that G is played repeatedly in discrete time periods $t = 0, 1, 2, \dots, T$. We denote $\pi_{t,i}(c_t, s_t^{-i})$ player i 's payoffs at time t , when they choose $c \in S_i$ (denoted c_t), assuming the others chose s^{-i} (denoted s_t^{-i}). Similarly, we use $BR_{t,i}$ to denote the set of best replies of player i at time t , when the opponents choose s_t^{-i} . The following subsections describe three revision protocols, by providing different possible probability distributions used by i at time $t+1$, contingent on i 's behavioral traits.

3.2. The 'Short-Sightedness' Protocol

The short-sightedness protocol is described by the following probability distribution:

$$p_{t+1,i}(y) = \mu \cdot (\pi_{t,i}(h_t, s_t^{-i}) - \pi_{t,i}(c_t, s_t^{-i})) / (m_i - 1) \cdot \Delta, \text{ for all } y \neq c, y \in S_i,$$

$$p_{t+1,i}(c) = 1 - \mu \cdot (\pi_{t,i}(h_t, s_t^{-i}) - \pi_{t,i}(c_t, s_t^{-i})) / \Delta,$$

$$\text{where } h_t \in BR_{t,i}, 0 < \mu \leq 1. \quad (11)$$

In words, unless player i chooses a best reply at time t , she may switch to any one among her alternative strategies with equal probability, one that is proportional to the ratio of the payoffs foregone at time t (the 'disappointment' at time t) over the maximum payoffs that could have been lost (the 'maximum disappointment' Δ). Thus, the greater the loss of payoffs at t (due to one's failure

of selecting a best reply), the higher the probability that i switches to some alternative strategy. Values of μ less than 1 indicate inertia, which increases as μ decreases². Agents conforming to (1) are clearly short-sighted, for they just take into account what has happened in the previous round only.

3.3. The ‘ n -Period Memory’ Protocol

In this revision protocol, players are depicted as more rigid or patient, as they switch only after experiencing n disappointing outcomes in a row (e.g., in rounds $t, t-1, \dots, t-n+1$). The probability distribution of this protocol is given by (2) below:

$$\begin{aligned} & \text{If } \Pi_{\tau=t-n+1}^t (\pi_{\tau,i}(h_\tau, s_\tau^{-i}) - \pi_{\tau,i}(c_\tau, s_\tau^{-i})) \neq 0, \\ & \text{then } p_{t+1,i}(y) = \sum_{\tau=t-n+1}^t (\pi_{\tau,i}(h_\tau, s_\tau^{-i}) - \pi_{\tau,i}(c_\tau, s_\tau^{-i})) / ((m_i - 1)n\Delta), \\ & \text{for all } y \neq c, y \in S_i, p_{t+1,i}(c) = 1 - \sum_{\tau=t-n+1}^t (\pi_{\tau,i}(h_\tau, s_\tau^{-i}) - \pi_{\tau,i}(c_\tau, s_\tau^{-i})) / n\Delta. \end{aligned} \quad (2)$$

$$\text{If } \Pi_{\tau=t-n+1}^t (\pi_{\tau,i}(h_\tau, s_\tau^{-i}) - \pi_{\tau,i}(c_\tau, s_\tau^{-i})) = 0, \text{ then } p_{t+1,i}(y) = 0,$$

for all $y \neq c, y \in S_i, p_{t+1,i}(c) = 1$, where $h_\tau \in BR_{\tau,i}$, and $\pi_{\tau,i}(h_\tau, s_\tau^{-i}) = \pi_{\tau,i}(c_\tau, s_\tau^{-i}) = 0$ when $\tau < 0$.

If $n = 1$, (2) collapses to (1) for the special case where $\mu = 1$; if $n > 1$, inertia comes into play with n reflecting the agent’s ‘resistance’ to switching following a string of disappointing outcomes. Note that protocol (2) implies that all past n rounds matter the same. This is plausible to the extent that n is quite small, and that agent i gets involved in the same interaction frequently enough. Discounting of older disappointing outcomes will be considered below.

3.4. The ‘Additive Disappointment’ Protocol

Let $v_i \in R^{T+1}$, $v_i = [v_{0,i} \ v_{1,i} \ \dots \ v_{T,i}]$ and define v_i as: $v_{0,i} = 1$; $v_{\tau,i} = v_{\tau-1,i} + 1$, $\tau \geq 1$, if $c_\tau = c_{\tau-1}$; $v_{\tau,i} = 1$, $\tau \geq 1$, if $c_\tau \neq c_{\tau-1}$. This vector effectively keeps track of when player i changes their strategy, and of how many rounds each strategy has lasted. For example, if player i chooses strategy $\beta \in S_i$ at $t = 0, t = 1, t = 2$, then strategy $\gamma \in S_i, \gamma \neq \beta$ at $t = 3$, and then strategy $\delta \in S_i, \delta \neq \gamma$ at $t = 4$ and $t = 5$, then the first six elements of v_i will be 1, 2, 3, 1, 1, 2. Hence, whenever we see that $v_{\tau,i}$ is equal to 1, we can know that a change of strategy happened at $t = \tau$, and that i had been choosing their previous strategy (the one they played until $t = \tau - 1$) for a number of rounds equal to $v_{\tau-1,i}$. We can now define the *additive disappointment* protocol by (3) below:

$$\begin{aligned} & \text{If } \sum_{\tau=t-v_{\tau,i}+1}^t \rho^{t-\tau} (\pi_{\tau,i}(h_\tau, s_\tau^{-i}) - \pi_{\tau,i}(c_\tau, s_\tau^{-i})) \geq Z, \text{ then } p_{t+1,i}(y) = 1 / (m_i - 1), \\ & \text{for all } y \neq c, y \in S_i, p_{t+1,i}(c) = 0. \end{aligned} \quad (3)$$

² As agent i is assumed as capable of acknowledging when they could have earned more payoffs, it would probably be more realistic to argue that, if i is to change their current strategy, then they are not going to select another strategy at random (as implied by the above distribution), but they shall choose a strategy belonging to the set of best replies to the opponents’ choice at t (i.e., the set $BR_{t,i}$). In the case of 2×2 games (that shall be studied here), it is obvious that this issue is not a concern. Obviously, for larger games, this protocol reflects a very weak form of learning, and should probably be modified in accordance to the sophistication one would wish to endow the individuals with.

If $\sum_{\tau=t-v_{t,i}+1}^t \rho^{t-\tau} \cdot (\pi_{\tau,i}(h_{\tau}, s_{\tau}^{-i}) - \pi_{\tau,i}(c_{\tau}, s_{\tau}^{-i})) < Z$, then $p_{t+1,i}(y) = 0$,
 for all $y \neq c$, $y \in S_i$, $p_{t+1,i}(c) = 1$, where $h_{\tau} \in BR_{\tau,i}$, $0 < \rho \leq 1$, and $Z \in R^+$.

The above reflects players who change their current strategy c at $t+1$ if and only if the amassed disappointment from the previous $v_{t,i}$ rounds (where c was played) surpasses a predefined threshold Z . This protocol describes players who abide with a strategy for a number of rounds until total disappointment exceeds some subjective ‘‘tolerance’’ level. In each new period, the disappointment received from the previous $v_{t,i}$ rounds is discounted at rate ρ . For simplicity, the population will be thought of as homogeneous in terms of parameters ρ and Z .

4. The Simulation

The software assumes a population of 1,000 agents and is amenable to symmetric 2×2 games³. The user specifies a number of iterations x . In each iteration or period, agents are randomly paired. The user may introduce a positive probability a with which an agent is, unbeknownst to her, matched against an ‘‘intruder’’; that is, an ‘outsider’ who does not belong to the original population and whose behavior has not evolved endogenously. To keep things simple, it is assumed that intruders act like automata, selecting their first strategy with probability b in every round and independently of their own past experience. In effect, they represent stochastic perturbations of the original population’s evolutionary process. Any values for a and b in $[0,1]$ are permitted (with 0.1 and 0.5 respectively being the default values). After the end of the simulation, each player will have participated in approximately $x \cdot (2 - a)/1,000$ rounds. With no intruders ($a = 0$), this number collapses to $x/500$.

At the outset, agents choose strategies with predefined (by the user) probabilities, with p_t denoting the fraction of the population that chooses the first strategy at time t . As we focus on symmetric games, this is also the probability that a random individual chooses their first strategy at t . The default value for p_0 is 0.5.

One crucial difference of our model with the canonical version of evolutionary game theory lies in that it makes no allusion to expected values, but rather, uses realized ones. This leads to a more plausible representation of the evolutionary course, consistent with the empirical assertion that someone’s current choices will reflect their tendency to avoid unpleasant past experiences due to perceived erroneous choices. The specific mechanics of the tendency to switch strategies following such ‘disappointment’ depends on the three behavioral types described in the previous section. The modifications of revision protocols (1), (2), and (3) necessary to introduce them into the software presented therein are straightforward. As only two players are randomly selected to participate in each round, the probability distributions (1), (2), and (3) are valid only for the periods where player i participates. Thus, periods $t = 0, 1, 2, \dots, T$, as used in these distributions, are no longer all the rounds of the game, but only those involving player i ⁴.

³ The software was written by the author in Microsoft Visual Basic 6.0. Random numbers are generated by use of the language’s Rnd and Randomize functions.

⁴ As this adjustment causes no ambiguity, we will simplify notation by not adding i subscripts to the time periods, as a more formal representation would require.

The following subsections present the modifications to protocols (1), (2), and (3), as implemented by the software program.

4.1. Short-Sightedness

We denote the currently used strategy with c , and the alternative one with y , while the opponent's current strategy is s^i . The time periods indicate the rounds where i is randomly selected to play the game, while parameter μ (which determines the maximum probability of a switch) has been hard-coded equal to 1/3. Protocol (1) is rewritten as follows:

$$\begin{aligned} p_{t+1,i}(y) &= (\pi_{t,i}(h_t, s_t^i) - \pi_{t,i}(c_t, s_t^i))/3\Delta, y \neq c, y \in S_i, \\ p_{t+1,i}(c) &= 1 - (\pi_{t,i}(h_t, s_t^i) - \pi_{t,i}(c_t, s_t^i))/3\Delta, \\ &\text{where } h_t \in BR_{t,i}. \end{aligned} \quad (1')$$

4.2. Three-Period Memory

In addition to its intuitive appeal, the introduction of behavioral inertia helps explain why the evolutionary process manages to escape its original state p_0 (see Subsection 4.5 for an example). Probability distribution (2') below is a special case of revision protocol (2) for 2×2 games and with $n = 3$. Here, switching to the other strategy is a possibility only when the current strategy has generated three disappointing outcomes in a row; an arbitrary, yet quite plausible assumption (supported also by popular phrases such as 'to err thrice...').

$$\begin{aligned} &\text{If } \prod_{\tau=t-2}^t (\pi_{\tau,i}(h_\tau, s_\tau^i) - \pi_{\tau,i}(c_\tau, s_\tau^i)) \neq 0, \\ &\text{then } p_{t+1,i}(y) = \sum_{\tau=t-2}^t (\pi_{\tau,i}(h_\tau, s_\tau^i) - \pi_{\tau,i}(c_\tau, s_\tau^i))/3\Delta, y \neq c, y \in S_i, \\ &p_{t+1,i}(c) = 1 - \sum_{\tau=t-2}^t (\pi_{\tau,i}(h_\tau, s_\tau^i) - \pi_{\tau,i}(c_\tau, s_\tau^i))/3\Delta. \end{aligned} \quad (2')$$

$$\begin{aligned} &\text{If } \prod_{\tau=t-2}^t (\pi_{\tau,i}(h_\tau, s_\tau^i) - \pi_{\tau,i}(c_\tau, s_\tau^i)) = 0 \text{ then } p_{t+1,i}(y) = 0, y \neq c, y \in S_i, p_{t+1,i}(c) = 1, \\ &\text{where } h_\tau \in BR_{\tau,i}, \text{ and } \pi_{\tau,i}(h_\tau, s_\tau^i) = \pi_{\tau,i}(c_\tau, s_\tau^i) = 0 \text{ when } \tau < 0. \end{aligned}$$

Three-period memory implies patient players who do not switch strategies at the first, or second setback. Change of current strategy c will not be considered if c has been the best reply in any one of the last three periods.

4.3. Additive Disappointment

Additive disappointment offers a less stylized variation of the previous protocol: agent i 's memory extends back to as many periods as the number of rounds that i has played c . The software assumes that disappointment in round $t-\tau$ is discounted at a 0.9^τ rate. The threshold value Z is fixed equal to two times the maximum disappointment level Δ . Hence, and given the discounting, i cannot switch to $y \neq c$ in fewer than three consecutive disappointing rounds.

$$\begin{aligned} &\text{If } \sum_{\tau=t-v_{i,j}+1}^t 0.9^{t-\tau} \cdot (\pi_{\tau,i}(h_\tau, s_\tau^i) - \pi_{\tau,i}(c_\tau, s_\tau^i)) \geq 2\Delta, \\ &\text{then } p_{t+1,i}(y) = 1, y \neq c, y \in S_i, p_{t+1,i}(c) = 0. \end{aligned} \quad (3')$$

$$\text{If } \sum_{\tau=t-v,i+1}^t 0.9^{t-\tau} \cdot (\pi_{\tau,i}(h_{\tau}, s_{\tau}^{-i}) - \pi_{\tau,i}(c_{\tau}, s_{\tau}^{-i})) < 2\Delta,$$

$$\text{then } p_{t+1,i}(y) = 0, y \neq c, y \in S_i, p_{t+1,i}(c) = 1,$$

$$\text{where } h_{\tau} \in BR_{\tau,i}.$$

Under this protocol, players switch strategies when they feel they have ‘had enough’ of the disappointment due to their strategy choice. The specified discount factor and disappointment threshold, while arbitrary, models, nicely, agents who adopt satisficing behavior in that they stick to some possibly suboptimal strategy until their tolerance level (factoring in the relative importance of more recent disappointment) is exceeded.

4.4. The Games

Table 1 shows the reference games to be used in the simulation:

Table 1. Five classic games.

	1	2		1	2		1	2		1	2		1	2		1	2		
1	3,3	0,5		1	2,2	0,0		1	2,2	0,0		1	-2,-2	2,0		1	3,3	0,2	
2	5,0	1,1		2	0,0	2,2		2	0,0	1,1		2	0,2	1,1		2	2,0	1,1	
Game 1:			Game 2:			Game 3:			Game 4:			Game 5:							
Prisoner's Dilemma			Coordination			Hi-Lo			Hawk-Dove			Stag-Hunt							

In each of these games, the ‘social welfare’ standpoint rejects certain outcomes (as inferior). Unfortunately, individually rational action often leads players to these very outcomes. In the ‘*Prisoner's Dilemma*’ and ‘*Hawk-Dove*’ games, for example, the socially desirable strategies cooperation and dovish behavior are trumped by defection and hawkish behavior (as a part of a mixed strategy) respectively, while, in ‘*Coordination*’, rationality alone cannot prevent coordination failure—and the same applies for ‘*Hi-Lo*’ and ‘*Stag-Hunt*’. Our interest is to see whether behavior evolving according to the above protocols leads to results that differ substantially from those suggested by standard game theory. The simulation results are described below.

4.5 Simulation Results: Case without Random Perturbations

In this subsection, we assume that there are no intruders ($a = 0$), and therefore, members of the population always interact between themselves. Interactions with intruders shall be seen as random ‘shocks’, and they will be presented in subsection 4.6.

4.5.1. Prisoner's Dilemma and Hawk-Dove

Unsurprisingly, in the ‘*Prisoner's Dilemma*’ defection dominates the entire population, no matter the behavioral type and the initial conditions. As in standard evolutionary game theory, defection always yields zero disappointment, while cooperation always yields positive disappointment; thus,

the agents shall eventually switch to defection, regardless of their behavioral code, insofar this disappointment is assumed to generate a tendency to defect, if one cooperates. Of course, things might have turned out otherwise if disappointment were to be defined differently. Here, disappointment is the feeling one gets when comparing one's actual payoff to what it would have been had one chosen a best reply strategy, the opponent's choice being a given. Naturally, if disappointment were to increase in proportion to one's share of foregone collective payoffs, then mutual cooperation would be a possibility, as long as players cared about the welfare of both participants as a group, and not just for their own performance in the game. One way to incorporate this consideration would typically be to change the payoffs of the game in order to reflect the increase in utility possibly derived from a mutually beneficial outcome and then rerun the simulation for the amended game. This would commonly give us a game with the strategic structure of '*Stag-Hunt*', the simulation results for which are presented below.

In '*Hawk-Dove*', short-sightedness causes evolution to converge to a state where around 41% choose the hawkish strategy *for any initial condition*; a level of aggression considerably higher to that predicted by standard evolutionary game theory, where the evolutionary stable equilibrium is $p = 1/3$. The aggression, in fact, grows further under the three-period memory protocol to, approximately, $p = 0.46$. The explanation for this last result is that the players' relative rigidity acts as an enabler for the existence of more 'hawks' in the population; under "three-period memory", players who behave aggressively need to be paired with other aggressive players in three consecutive rounds before changing their strategy to 'dove', and this persistence tolerates more 'hawks' in the aggregate than in the case of myopic players who switch to the other strategy more readily. Finally, under additive disappointment, the system converges to approximately $p = 0.41$ (same level approximately as under short-sightedness). This rate of aggression decreases dramatically as the payoff consequences of a (Hawk–Hawk) outcome become more disastrous: if, for example, the payoffs of the (Hawk–Hawk) outcome are changed from $(-2, -2)$ to $(-4, -4)$, then the new equilibrium becomes $p = 0.25$ (N.B., the other protocols give $p = 0.33$ (short-sightedness) and $p = 0.43$ (three-period memory)). This decrease in 'hawks' is explained on the grounds that, as the disappointment caused by a conflictual outcome became four times greater than the disappointment generated by mutual acquiescence (Dove–Dove) on the occasion that 'hawks' cross paths, it is much more probable that the threshold value determining whether a change of strategy will happen or not is surpassed (as opposed to the case where two 'doves' meet).

4.5.2. Coordination, Hi-Lo, and Stag Hunt

In these three variants of the coordination problem, standard evolutionary game theory admits each game's pure Nash equilibria as evolutionarily stable. Which of the two obtains depends on the initial conditions: In '*Coordination*' and '*Stag Hunt*', the top left (bottom right) equilibrium, $p = 1$ ($p = 0$), will emerge if, at the outset, more (less) than 50% of the population chose the first strategy (*i.e.*, if $p_0 > 1/2$). In '*Hi-Lo*', the top left equilibrium requires a smaller initial critical mass in order to dominate ($p_0 > 1/3$).

In contrast, under the *short-sightedness* protocol presented in section 4.1, above, convergence to one of the two equilibria is not as straightforward. In the pure '*Coordination*' game, while the system

has no tendency to leave these equilibria once there (as no individual ever experiences disappointment at these equilibria), for any other initial condition ($0 < p_0 < 1$) the system will perform a random walk, oscillating back and forth in each round with equal probabilities, to only rest if one of the two equilibria is accidentally reached. The explanation for this lies in that both players have the same probability of changing their strategy when coordination fails; therefore, the probability for p_t to increase in some round t is the same as the probability for p_t to decrease (and equal to $2/9 = \mu(1 - \mu)$); and if both players switch their strategies, or if neither does, then p_t remains unchanged.

It follows that under the short-sightedness protocol, agents caught in a pure ‘*Coordination*’ problem do a poor job of achieving coordination, resembling the embarrassing situation of two people trying to avoid collision when walking toward one another in some corridor. This result also shows why *some* inertia may be useful: for if μ were equal to 1, then, in any instance where one player would choose the first strategy and the other player would choose the second strategy, then both would switch to their other strategy with probability 1, and, hence, $p_t = p_0$ for all t .

The situation is different when the short-sightedness scenario is used in the context of ‘*Hi-Lo*’. In that game, the fact that the equilibria are Pareto-ranked helps the system converge to the optimal equilibrium $p = 1$, for any initial condition. The only exception is, naturally, the extreme case where $p_0 = 0$. In short, the efficient outcome is attained as long as a minor fraction of individuals is opting, initially, for the first strategy, because an instance of non-coordination in this game brings greater disappointment to the player who chose the second strategy; hence, the probability with which the player who chose the second strategy will switch to the first strategy is greater than the probability with which the player who chose strategy 1 will switch to strategy 2. This explains the convergence to the efficient outcome from any state except for the one where (nearly) all players choose the second strategy.

Turning now to the *three-period memory* protocol, in pure ‘*Coordination*’, the results coincide with those of standard evolutionary game theory (if $p_0 = 0.5$, then the system converges to either $p = 0$ or $p = 1$ with equal probabilities; if $p_0 < (>) 0.5$, it tends to $p = 0(1)$). Here, the relative rigidity of the players is what enables them to arrive at an equilibrium, for switching to the other strategy may only happen after three consecutive disappointing outcomes, which makes a change of strategy more probable when fewer people are choosing it. In juxtaposition, applying the three-period memory protocol to ‘*Hi-Lo*’ leads to the result that the necessary initial condition for emergence of the efficient equilibrium is approximately $p_0 > 0.42$. When the equilibria are Pareto-ranked, and unlike the situation in which agents are short-sighted, the evolution of the optimal equilibrium requires a lot more initial adherents (*i.e.*, a much higher p_0). Clearly, the players’ relatively high inertia may inhibit the evolution of the Pareto optimal outcome since, if a critical mass of a least 42% of players opting for the Pareto optimal strategy is not present from the outset, it is sufficiently likely that those who choose it will experience three disappointing results in a row, causing them to switch to the strategy that corresponds to the suboptimal equilibrium.

On the other hand, under the *additive disappointment* protocol, coordination at the optimal equilibrium of ‘*Hi-Lo*’ is more probable (though still less so than in the short-sightedness protocol case) as the efficient equilibrium’s catchment basin is approximately $p_0 > 0.17$. Interestingly, when the suboptimal strategy (strategy 2) carries less risk in case of coordination failure, as in the ‘*Stag*

Hunt game, the way in which disappointment affects the players makes little difference: under both the three-period memory and the additive disappointment protocols, efficiency is guaranteed as long as $p_0 > 1/2$, and condemned if $p_0 < 1/2$.

It is worth noticing that the analysis is sensitive to the relative attractiveness of the efficient equilibrium (just as the relative unattractiveness of the conflictual outcome made a crucial difference in *Hawk-Dove*). If, for example, we change the payoffs of the efficient outcome of *Stag-Hunt* from (3,3) to (4,4), then we see that, under short-sightedness, the system converges to the efficient equilibrium from *any* initial state (except for the case where $p_0 = 0$). However, this can also work the other way round: if we make the inferior outcome less unattractive, then we get the same effect in reverse. For example, if we change the payoffs of the sub-optimal outcome of *Stag-Hunt* from (1,1) to (1.5,1.5), then the system shall always converge to $p = 0$ (unless $p_0 = 1$). In this last example, the three-period memory and the additive disappointment protocols lead to the efficient equilibrium as long as $p_0 > 0.55$ and $p_0 > 0.74$, respectively.

Table 3 summarizes the above results for the games of Table 1 and the amended games featured in Table 2.

Table 2. Games 4 and 5, amended.

	1	2		1	2		1	2
1	-4,-4	2,0	1	4,4	0,2	1	3,3	0,2
2	0,2	1,1	2	2,0	1,1	2	2,0	1.5,1.5
	Game 6:			Game 7:			Game 8:	
	Hawk-Dove #2			Stag-Hunt #2			Stag-Hunt #3	

It is now clear that the agents' behavioral type is a crucial determinant of the evolutionary path. The players' attitude to disappointing resolutions may not make a difference in games with unique dominant strategy equilibria, like the *Prisoner's Dilemma*, but it does so in games featuring multiple Nash/evolutionary equilibria, e.g., *Hawk-Dove*, pure *Coordination*, *Hi-Lo*, or *Stag-Hunt*.

Our first insight is that, while short-sighted players perform poorly when coordinating on equally desirable equilibria (e.g., pure *Coordination*), they may be more adept than players with longer memories at sidestepping paths that railroad them toward inefficient equilibria (as in *Hi-Lo*). The same may be also true for coordination type games where Pareto-efficiency and an aversion to the worst outcome may pull players in different directions, e.g., *Stag-Hunt*. On the other hand, as the relative benefits from the efficient outcome decrease, the agents' myopia may have the opposite effect (recall *Stag-Hunt #3*).

Table 3. Simulation results in the case without random perturbations.

	Short-Sightedness	Three-Period Memory	Additive Disappointment
Prisoner's Dilemma		$p = 0$ from any initial state	
Coordination	Randomness, possible convergence to $p = 0$ or $p = 1$	$p = 1$, if $p > 1/2$ at $t = 0$ $p = 0$, if $p < 1/2$ at $t = 0$	$p = 1$, if $p > 1/2$ at $t = 0$ $p = 0$, if $p < 1/2$ at $t = 0$
Hi-Lo	$p = 1$, if $p > 0$ at $t = 0$ $p = 0$, if $p = 0$ at $t = 0$	$p = 1$, if $p > 0.42$ at $t = 0$ $p = 0$, if $p < 0.42$ at $t = 0$	$p = 1$, if $p > 0.17$ at $t = 0$ $p = 0$, if $p < 0.17$ at $t = 0$
Hawk-Dove	$p \approx 0.41$	$p \approx 0.46$	$p \approx 0.41$
Hawk-Dove #2	$p \approx 0.33$	$p \approx 0.43$	$p \approx 0.25$
Stag-Hunt		Same results as ' <i>Coordination</i> '	
Stag-Hunt #2		Same results as ' <i>Hi-Lo</i> '	
Stag-Hunt #3	$p = 0$, if $p < 1$ at $t = 0$ $p = 1$, if $p = 1$ at $t = 0$	$p = 1$, if $p > 0.55$ at $t = 0$ $p = 0$, if $p < 0.55$ at $t = 0$	$p = 1$, if $p > 0.74$ at $t = 0$ $p = 0$, if $p < 0.74$ at $t = 0$

Agents described by the additive disappointment protocol fare better in antagonistic interactions; their attitude of sticking to a strategy (unless the amassed disappointment from previous rounds surpasses some threshold) is bound to turn them into a peaceful group of people, not necessarily because they favor peace, but because they are contended more easily and have no incentive to strive for more, at peace's expense. Moreover, these agents perform remarkably well in '*Hi-Lo*', '*Stag-Hunt*', and '*Stag-Hunt #2*' (albeit worse than short-sighted agents), but not so well in '*Stag-Hunt #3*', where the catchment area of the efficient equilibrium is relatively small.

The three-period memory protocol credits the agents with some level of sophistication. Their elevated inertia does not seem to be in their favor in several cases, especially in '*Hawk-Dove*', where the resulting aggression is too high, and in '*Hi-Lo*' or '*Stag-Hunt #2*', where the basin of attraction of the efficient equilibrium is smaller than the other behavioral types; however, their sense of caution pays off in '*Stag-Hunt #3*', where they are ultimately driven to the optimal outcome even for relatively low initial p_0 values. These players are sometimes not flexible enough to let the evolutionary course work in their favor, but this very rigidity is what may protect them against possibly unpleasant situations (such as being attracted by the sub-optimal equilibrium in '*Stag-Hunt #3*' from any initial state except $p_0 = 1$, as happens under short-sightedness).


4.6 Simulation Results: Case with Random Perturbations

4.6.1. Short-Sightedness with Stochastic Perturbations

We have already noticed how short-sighted players may be heavily influenced by minor perturbations. To explore this further, the simulation software has been augmented potentially to include ‘intruders’. The latter interact with our population members with probability a (in each iteration) and choose the first strategy with probability b . Their presence may be interpreted either as a random error or a deliberate intervention—perhaps from a third party aiming to help the original population arrive at a desired equilibrium. These perturbations might as well be considered as noise that sometimes enables (as shall be seen below) the dynamics to leave a catchment area and enter another basin of attraction.

To illustrate, let us consider ‘*Hi-Lo*’ under short-sightedness. In the case without perturbations, we saw that the efficient equilibrium is threatened only if the whole population is stuck, at the very beginning, in the suboptimal equilibrium. Naturally, the introduction of only a few intruders is enough to guarantee convergence to the efficient equilibrium $p = 1$. Figure 1 demonstrates this under the assumption that $p_0 = 0$, $a = 0.01$, and $b = 0.1$: as time goes by (horizontal axis), the number of individuals (out of 1000) who choose strategy 1 grows inexorably (the vertical axis depicts the number of individuals who choose their first strategy or, equivalently, p_t multiplied by 1000). Convergence to the efficient outcome took, in this simulation, around 80 games per person (less than 40,000 iterations in total).


Figure 1. ‘*Hi-Lo*’ when all players are initially ‘stuck’ in the inefficient outcome. The introduction of few intruders sets them on a course to the efficient outcome.



We now turn to ‘*Stag-Hunt*’, where the evolutionary path may take the population to one of the two available equilibria. Without stochastic perturbations, short-sightedness threatened to put the population in an endless drift (see previous section). Typically, if $p_0 = 0.5$, our (non-stochastic) simulation took more than one million iterations for the system to hit one of the two absorbing barriers. Naturally, the closer p_t is to one of the two barriers/equilibria, the more probable convergence is to that equilibrium point. However, the addition of intruders can change this. Consider the case where $p_0 = 0.1$. In Figure 2, Series 1 shows the results of a simulation without intruders: predictably, the proximity of the system’s initial condition to the inefficient outcome causes the population to converge toward it quite quickly. However, the addition of a small number of intruders

(1% of the population, *i.e.*, $a = 0.01$, who always play the first strategy, *i.e.*, $b = 1$) gives rise to Series 2 and, thus, to a drastically different path. The intruders' presence becomes the catalyst which creates what could be called as 'optimism' within the group, and ultimately drives it towards the optimal equilibrium.

Figure 2. 'Stag-Hunt' when 90% of players are, initially, drawn to the inefficient outcome (Series 1). The introduction of few intruders sets them on a course to the efficient outcome (Series 2).



Notwithstanding the obvious merits of short-sightedness, the implied impatience of the agents and the ease with which they switch to the other strategy may not always be a virtue. Figure 3 demonstrates the point: Here, 80% of the population is drawn to the efficient outcome and yet the presence of a similar number of intruders (as in Figure 2; namely $a = 0.01$ and $b = 0$) causes the evolutionary path to take the population straight into the arms of the inefficient outcome.

Our analysis in the case of no intruders in the previous section suggested that, in games of the 'Stag-Hunt' structure, short-sighted players may be easy to manipulate, depending on the relative gains from achieving a Pareto superior outcome. Under shortsightedness, players are assumed to experience equal disappointment from an instance of non-coordination in 'Stag-Hunt', regardless of whether they chose the first or the second strategy. In 'Stag-Hunt #3', however, the player who chooses the first strategy receives more disappointment than the player who chooses the second strategy, and hence, the intervention for emergence of the efficient equilibrium needs to be more drastic. Figure 4 presents a relevant simulation, with the initial condition $p_0 = 0.5$. Without intruders, the evolutionary course would have taken the population to the sub-optimal outcome (Series 1). However, a sizeable population of intruders may avert this: with $a = 0.5$ and $b = 1$ (that is, if all agents have a 1 in 2 chance of meeting an intruder who always selects the first strategy), the efficient outcome is guaranteed (Series 2). We notice that, on the one hand, the efficient equilibrium is attained, but, on the other hand, we can no longer speak of a minor perturbation or an uncalculated error: the intervention here has to be quite radical.

Figure 3. ‘*Stag-Hunt*’ when 80% of players are, initially, drawn to the efficient outcome. The introduction of few intruders sets them on a course to the inefficient outcome.





Figure 4. ‘*Stag-Hunt* #3’ when 50% of players are, initially, drawn to the efficient outcome. When 50% of them are intruders the efficient outcome is guaranteed (Series 2).



4.6.2. Three-Period Memory with Stochastic Perturbations

The three-period memory behavioral code is generally resistant to shocks. While a minor shock may have dramatic effects under short-sightedness, the same does not apply when agents are more patient, even when the perturbation is far from discrete.

In ‘*Hawk-Dove*’, we found that a three-period memory protocol with no random perturbations increased the players’ observed aggression. When a significant probability of meeting a hawkish intruder is introduced, one might be excused to expect a considerable drop in aggression. But that is not what we find in our simulation results. Figure 5 compares a simulation when there are no intruders (Series 1) with one in which there is a 33% probability ($a = 0.33$) of meeting an intruder who always chooses ‘Hawk’ ($b = 1$). The initial condition is $p_0 = 0.5$. We find that, while the percentage of aggressive players indeed decreases, the effect is rather minor (the difference of the two series is less than 10%, which is insignificant *viz.* a perturbation involving 1 in 3 games played by every person).

Turning to ‘*Hi-Lo*’, and recalling that the three-period memory protocol led to suboptimal results for the population as a whole, an infusion of intruders may make the necessary difference as long as their number is high enough. To give one example, in Figure 6 we set $p_0 = 0.35$. Without intruders, as we saw in previous sections, the system will rest at the inefficient equilibrium ($p = 0$), in contrast to the short-sightedness and the additive disappointment protocols, where $p = 1$ is the equilibrium. Nothing changes here when the proportion of intruders is small. Series 1 of Figure 6 demonstrates this amply. However, when the proportion of intruders rises to approximately 15%, a different path becomes possible, one that take the population to the optimal outcome. More precisely, for 10 different simulations of the same scenario with $p_0 = 0.35$, $a = 0.15$, and $b = 1$, there were six instances of convergence to the efficient equilibrium. Series 2 shows one of these instances (N.B., the smaller p_0 , the greater the value of a necessary for the system to converge to $p = 1$).

Figure 5. ‘*Hawk-Dove*’ with $p_0 = 0.5$. Series 1: $a = 0$, Series 2: $a = 0.33$, $b = 1$.





Figure 6. ‘*Hi-Lo*’ with $p_0 = 0.35$. Series 1: $a = 0.1$, $b = 1$. Series 2: $a = 0.15$, $b = 1$.



Meanwhile, in ‘*Stag-Hunt*’, the inefficient result will not be avoided even in the presence of a sizeable population of intruders. Figure 7 shows that even if, say, 40% of the population are drawn initially to the ‘good’ strategy *and* there is a probability of 15% of meeting with an intruder who also plays the ‘good’ strategy, the efficient equilibrium (that requires players to choose their ‘good’ strategies) will not eventuate. Naturally, in some circumstances, such rigidity may turn out to be in

the players' favor. In 'Stag-Hunt #3', with no intruders, we have already observed how short-sighted agents are attracted to the sub-optimal equilibrium (as a player who chooses the second strategy has a smaller probability of changing their strategy than a player who chooses the first strategy when the outcome of some round is (1,2) or (2,1)). In the case of the three-period memory protocol, the sub-optimal outcome has $p_0 < 0.55$ as its catchment area, and when p_0 is slightly greater than that, the efficient equilibrium is not threatened, not even when there is a 10% percent probability of meeting an intruder who always selects the second strategy ($a = 0.1, b = 0$). Figure 8 offers a relevant simulation with $p_0 = 0.6$. Series 1 emerges when there are no intruders, while Series 2 illustrates the scenario $a = 0.1, b = 0$. Note how the efficient equilibrium is reached either way, albeit at different speeds depending on the preponderance of intruders.

Figure 7. 'Stag-Hunt' with $p_0 = 0.4, a = 0.15, b = 1$.





Figure 8. 'Stag-Hunt #3' with $p_0 = 0.6$. Series 1: $a = 0$. Series 2: $a = 0.1, b = 0$.




4.6.3. Additive Disappointment with Stochastic Perturbations

Our additive disappointment protocol stands as some kind of middle ground between the impatience of the short-sighted players and the inertia of agents behaving under the three-period memory protocol. This section concludes with several instructive scenarios based on the disappointment protocols.

Under additive disappointment, Figure 9 shows that the presence of intruders lowers the population's aggression rate in 'Hawk-Dove', although the effect is not distinctly large. In 'Hi-Lo', the efficient outcome seems to have a surprisingly big basin of attraction ($p_0 > 0.17$). Even if $p_0 < 0.17$, a small perturbation is enough to avert convergence to the sub-optimal equilibrium (see Figure 10 which suggests that players who conform to the additive disappointment protocol are more prone to external influences than agents acting upon three-period memory, albeit they are not as impulsive as the short-sighted players). In 'Stag-Hunt', we notice a similar effect: Figure 11 shows that lack of intruders means convergence to the socially lesser equilibrium (Series 1, with $p_0 = 0.1$ and $a = 0$), whereas an infusion of 10% intruders suffices to energize a path like that of Series 2. Once more, we find that the population needs only a mild external influence to avoid unpleasant consequences but shows less flexibility when compared to short-sighted players.

Figure 9. 'Hawk-Dove' with $p_0 = 0.5$. Series 1: $a = 0$, Series 2: $a = 0.1$, $b = 1$.



The relative inertia of the additive disappointment protocol is also illustrated in Figure 12 which describes the evolutionary course for 'Stag-Hunt #3' with $p_0 = 0.5$, $a = 0.25$, and $b = 1$. Even though the proportion of intruders is quite large (25%), this is not enough to favor the optimal equilibrium. However, by the same token, the population does not converge to the suboptimal equilibrium either: instead, the system seems to wander around a non-equilibrium state in the proximity of $p = 0.15$. In that state, it is *as if* the disappointment received from instances of coordination failure is too weak to generate behavioral changes. Thus, some behavioral equilibrium (akin to satisficing) emerges at which players experience too little of an urge to switch to the other strategy, feeling that their chosen

behavior works well enough for them. The real benefits from switching to the optimal strategy are simply not large enough at the level of the individual.

Figure 10. ‘*Hi-Lo*’ with $p_0 = 0.1$. Series 1: $a = 0$, Series 2: $a = 0.05$, $b = 1$.




Figure 11. ‘*Stag-Hunt*’ with $p_0 = 0.1$. Series 1: $a = 0$, Series 2: $a = 0.1$, $b = 1$.





Figure 12. ‘*Stag-Hunt #3*’ with $p_0 = 0.5$, $a = 0.25$, $b = 1$.



Tables 4–6 summarize the results of the simulations presented in Subsections 4.6.1 to 4.6.3:

Table 4. Simulation results for short-sightedness protocol with stochastic perturbations.

Protocol: Short-Sightedness			
Game	Initial Condition	Convergence	Comments
Hi-Lo	$p = 0, a = 0.01, b = 0.1$	$p = 1$	A minor perturbation is enough for convergence to $p = 1$.
Stag-Hunt	$p = 0.1, a = 0.01, b = 1$	$p = 1$	If $a = 0$ (no perturbations) we have convergence at $p = 0$.
Stag-Hunt	$p = 0.8, a = 0.01, b = 0$	$p = 0$	If $a = 0$ (no perturbations) we have convergence at $p = 1$.
Stag-Hunt #3	$p = 0.5, a = 0.5, b = 1$	$p = 1$	If $a = 0$ (no perturbations) we have convergence at $p = 0$.

Table 5. Simulation results for three-period memory protocol with stochastic perturbations.

Protocol: Three-Period Memory			
Game	Initial condition	Convergence	Comments
Hawk-Dove	$p = 0.5, a = 0.33, b = 1$	$p \approx 0.38$	Convergence at $p \approx 0.47$ in the absence of perturbations.
Hi-Lo	$p = 0.35, a = 0.1, b = 1$	$p = 0$	
	$p = 0.35, a = 0.15, b = 1$	$p = 1$ (6 out of 10 runs)	
Stag-Hunt	$p = 0.4, a = 0.15, b = 1$	$p = 0$	Convergence at $p = 1$ under short-sightedness.
Stag-Hunt #3	$p = 0.6, a = 0.1, b = 0$	$p = 1$	If $a = 0$ convergence at $p = 1$ is quicker.

Table 6. Simulation results for additive disappointment protocol with stochastic perturbations.

Protocol: Additive disappointment			
Game	Initial Condition	Convergence	Comments
Hawk-Dove	$p = 0.5, a = 0.1, b = 1$	$p \approx 0.39$	Convergence at $p \approx 0.41$ in the absence of perturbations.
Hi-Lo	$p = 0.1, a = 0.05, b = 1$	$p = 1$	If $a = 0$ (no perturbations) we have convergence at $p = 0$.
Stag-Hunt	$p = 0.1, a = 0.1, b = 1$	$p = 1$	If $a = 0$ (no perturbations) we have convergence at $p = 0$.
Stag-Hunt #3	$p = 0.5, a = 0.25, b = 1$	$p \approx 0.15$	

5. Discussion and Conclusions

The simulations presented in the previous section allow for comparisons across the three behavioral ‘types’ modeled in Section 3 as (1’), (2’), and (3’). None of these ‘types’ is ‘best’, in the sense of boosting either individual or social welfare. Each type may perform better in one game and then worse in another. Therefore, in the hypothetical case in which agents have a *choice* as to *their* ‘type’, it might be optimal (insofar that this is possible) to adopt different ‘types’ depending on the interaction. For instance, agents may react to disappointment differently in interactions of an antagonistic nature (e.g., ‘*Hawk-Dove*’) from the way they react in cases of coordination failure (e.g., ‘*Hi-Lo*’ or ‘*Stag-Hunt*’). They may be more rigid in, say, ‘*Hawk-Dove*’ (possibly opting for the additive disappointment protocol) than in pure ‘*Coordination*’, where they may feel more relaxed and conform to the short-sightedness protocol. In fact, the simulation results illustrate that this would indeed be an advantageous tactic, given that none of the examined behavioral codes consistently

outperforms the others. Naturally, an interesting extension of this conclusion would be to provide a formal (and quantitative) definition of what constitutes a desired outcome in a game, so that one would be able to calculate the deviation of a specific rule (*i.e.*, behavioral code) from what is thought of as “best”.

The paper also illuminated the central role of stochastic perturbations in the determination of the relative social welfare effects of the different behavioral ‘types’. Short-sighted agents were shown to be highly sensitive to random perturbations and more likely to benefit from a benevolent third party (a social planner, perhaps?) who directs such shocks in a bid to steering the population to the desired equilibrium. On the other hand, if the planner’s intentions are not benign, the population risks being led to a sub-optimal outcome just as easily. In this sense, the three-period memory protocol shields a population from malevolent outside interventions at the expense of reducing the effectiveness of social policy that would, otherwise, have yielded effortless increases in social welfare.

Additive disappointment combines elements from both the short-sightedness and the three-period memory protocols, but its main theoretical disadvantage is that the system seems too sensitive to the choice of two exogenous parameters (the discount factor ρ and the threshold value Z). The simulation reported here ($\rho = 0.9$ and $Z = 2\Delta$) implies individuals with a moderate threshold of tolerance, in the sense that it is neither too high to prohibit strategy switches, nor too low to permit too much flexibility. In games featuring multiple evolutionary equilibria, a population of these ‘types’ may drift somewhere in-between the two equilibria (recall Figure 12). This is consistent with a novel type of behavioral equilibrium which does not correspond to any of the game’s evolutionary equilibria. Such a state of behavioral rest is more likely to occur in some form of coordination problem (e.g., ‘*Stag-Hunt*’, ‘*Hi-Lo*’), the result being a mixture of behaviors that, while stable, does not correspond to a mixed strategy equilibrium (in the traditional game theoretical sense). To give one real life example, when one observes that QWERTY and DVORAK typewriter keyboards are both still in use, then it is conceivable that this is a behavioral equilibrium state of the type simulated here.

On a similar note, for players acting under additive disappointment, there can be combinations of ρ and Z which yield positive cooperation rates in the ‘*Prisoner’s Dilemma*’ or unusually low aggression rates (or even zero) in ‘*Hawk-Dove*’. This kind of result is consistent with the observation of a stable proportion of ‘cooperators’ in the ‘*Prisoner’s Dilemma*’ (as confirmed by virtually all related experimental work), or players of ‘*Hawk-Dove*’ who bypass opportunities to behave aggressively when their opponents are acting dovishly. The analytical interpretation in this paper is that these players do not find the net benefits from switching to the other strategy high enough to motivate a change in their behavior. While the explanation of why that might be so might lie on bounded rationality considerations, it may also have its roots in the psychological character or social norms pertaining to the agents; e.g., perceptions of fairness or the intrinsic value of cooperation.

Finally, a critical note: This paper has confined its attention to homogeneous populations comprising agents who subscribe exclusively to one of the three revision protocols. A more realistic analysis would allow not only for the coexistence of these protocols in the same population but also for heterogeneity within a single protocol (*i.e.*, agents who, while adopting the additive disappointment protocol, feature different values for parameters ρ and Z). Future research along these

lines promises to throw important new light on the manner in which learning processes allow populations to achieve greater social and individual success in the games they play.

Acknowledgements

I am indebted to Yanis Varoufakis for his valuable comments. I would also like to thank three anonymous reviewers for their helpful suggestions.

Conflicts of Interest

The author declares no conflict of interest.

References

1. Sugden, R. The evolutionary turn in game theory. *J. Econ. Methodol.* **2001**, *1*, 113–130.
2. Rubinstein, A. *Modeling Bounded Rationality*; MIT Press: Cambridge, MA, USA, 1998.
3. Maynard Smith, J.; Price G.R. The logic of animal conflict. *Nature* **1973**, *246*, 15–18.
4. Maynard Smith, J. The theory of games and the evolution of animal conflicts. *J. Theor. Biol.* **1974**, *47*, 209–221.
5. Lewontin, R.C. Evolution and the theory of games. *J. Theor. Biol.* **1961**, *1*, 382–403.
6. Taylor, P.; Jonker L. Evolutionary stable strategies and game dynamics. *Math. Biosci.* **1978**, *40*, 145–156.
7. Mailath, G.J. Do people play Nash equilibrium? Lessons from evolutionary game theory. *J. Econ. Lit.* **1998**, *36*, 1347–1374.
8. Weibull, J.W. What have we learned from evolutionary game theory so far? Research Institute of Industrial Economics IUI, 1998, Working Paper Series 487. Available online: <http://swopec.hhs.se/iuiwop/papers/iuiwop0487.pdf> (accessed on 1 November 2013).
9. Samuelson, L. Evolution and game theory. *J. Econ. Perspect.* **2002**, *16*, 47–66.
10. Hofbauer, J.; Sigmund K. Evolutionary game dynamics. *B. Am. Math. Soc.* **2003**, *40*, 479–519.
11. Friedman, D. On economic applications of evolutionary game theory. *J. Evol. Econ.* **1998**, *8*, 15–43.
12. Schlag, K.H. Why imitate, and if so, how? A bounded rational approach to multi-armed bandits. *J. Econ. Theory* **1998**, *78*, 130–156.
13. Cross, J. A stochastic learning model of economic behavior. *Q. J. Econ.* **1973**, *87*, 239–266.
14. Borgers, T.; Sarin R. Learning through reinforcement and replicator dynamics. *J. Econ. Theory* **1997**, *77*, 1–14.
15. Erev, I.; Roth, A.E. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *Am. Econ. Rev.* **1998**, *88*, 848–881.
16. Fudenberg, D.; Kreps D. Learning mixed equilibria. *Game Econ. Behav.* **1993**, *5*, 320–367.
17. Fudenberg, D.; Levine D.K. *The Theory of Learning in Games*; MIT Press: Cambridge, MA, USA, 1998.
18. Cressman, R. *Evolutionary Dynamics and Extensive Form Games*; MIT Press: Cambridge, MA, USA, 2003.

19. Sandholm, W.H. *Population Games and Evolutionary Dynamics*; MIT Press: Cambridge, MA, USA, 2010.
20. Foster, D.; Young H.P. Stochastic evolutionary game dynamics. *Theor. Popul. Biol.* **1990**, *38*, 219–232.
21. Kandori, M.; Mailath G.J.; Rob R. Learning, mutation, and long run equilibria in games. *Econometrica* **1993**, *61*, 29–56.
22. Young, H.P. The evolution of conventions. *Econometrica* **1993**, *61*, 57–84.
23. Binmore, K.; Samuelson, L.; Vaughan, R. Musical chairs: Modeling noisy evolution. *Game Econ. Behav.* **1995**, *11*, 1–35.
24. Young, H.P. *Individual Strategy and Social Structure*; Princeton University Press, Princeton, NJ, USA, 1998.
25. Loomes, G.; Sugden, R. Regret theory: An alternative theory of rational choice under uncertainty. *Econ. J.* **1982**, *92*, 805–824.
26. Gul, F. A theory of disappointment aversion. *Econometrica* **1991**, *59*, 667–686.
27. Braun, M.; Muermann, A. The impact of regret on the demand of insurance. *J. Risk Insur.* **2004**, *71*, 737–767.
28. Irons, B.; Hepburn, C.J. Regret theory and the tyranny of choice. *Econ. Rec.* **2007**, *83*, 191–203.
29. Laciara, C.E.; Weber, E.U. Correcting expected utility for comparisons between alternative outcomes. *Journal of Risk and Uncertainty* **2008**, *36*, 1–17.
30. Grant, S.; Atsushi, K.; Polak, B. Different notions of disappointment aversion. *J. Econ. Lit.* **2001**, *81*, 203–208.
31. Hart, S.; Mas-Colell A. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* **2000**, *68*, 1127–1150.
32. Macy, M.W.; Flache, A. Learning dynamics in social dilemmas. *P. Natl. A. Sci.* **2002**, *99*, 7229–7236.
33. Hodgson, G.M.; Knudsen T. The complex evolution of a simple traffic convention: the functions and implications of habit. *J. Econ. Behav. Organ.* **2004**, *54*, 19–47.
34. Radax, W.; Wäckerle, M.; Hanappi, H. From agents to large actors and back; Formalized story-telling of emergence and exit in political economy. 2009. Available online: http://publik.tuwien.ac.at/files/PubDat_177962.pdf (accessed on 1 November 2013).
35. Heinrich, T.; Schwardt, H. Institutional inertia and institutional change in an expanding normal-form game. *Games* **2013**, *4*, 398–425.
36. Weibull, J.W. *Evolutionary Game Theory*; MIT Press, Cambridge, MA, USA, 1995.

Schelling, von Neumann, and the Event that Didn't Occur

Alexander J. Field

Abstract: Thomas Schelling was recognized by the Nobel Prize committee as a pioneer in the application of game theory and rational choice analysis to problems of politics and international relations. However, although he makes frequent references in his writings to this approach, his main explorations and insights depend upon and require acknowledgment of its limitations. One of his principal concerns was how a country could engage in successful deterrence. If the behavioral assumptions that commonly underpin game theory are taken seriously and applied consistently, however, nuclear adversaries are almost certain to engage in devastating conflict, as John von Neumann forcefully asserted. The history of the last half century falsified von Neumann's prediction, and the "event that didn't occur" formed the subject of Schelling's Nobel lecture. The answer to the question "why?" is the central concern of this paper.

Reprinted from Special Issue: Aspects of Game Theory and Institutional Economics, *Games*.
Cite as: Field, A.J. Schelling, von Neumann, and the Event that Didn't Occur. *Games* **2014**, *5*, 53–89.

1. Introduction

Thomas Schelling is widely thought of, and was recognized by the Nobel Prize committee as a pioneer in the application of game theory and rational choice analysis to problems of politics and international relations. Much of the popularity of his work and other analysis in this vein stemmed from the perception that it contributed to the development and application of new "tools" for understanding and analyzing social phenomena. Following the prize award, the economics journalist David Warsh described him as "the pioneering strategist who made game theory serve everyday economics for thirty years" [1, p. 7].

However, although Schelling makes frequent references in his writings to rational choice and game theory, his analysis of deterrence¹ is based on assumptions about human behavior and logic which, although useful in thinking practically about strategic policy, are at variance with those commonly adduced by game theorists, at least those specializing in its non-cooperative variant.² In areas especially relevant for strategy and conflict, game theory leads to behavioral predictions which are simply not borne out in the laboratory or, as will be apparent, in the real world.

In the Prisoner's Dilemma played once, for example, the Nash prediction is unambiguous: no cooperation. Defection is the strictly dominant strategy. Experimental evidence, however, provides

¹ Deterrence most commonly brings to mind the prevention of attacks on one's own territory or that of close allies. But it can, more aggressively, be used in furtherance of other foreign policy aims. Schelling was interested in both defensive deterrence and its more aggressive forms, and the role that nuclear arms might play in either.

² Non-cooperative theory studies interactions in which players are not allowed to make binding commitments among themselves. Cooperative theory allows such agreements, without specifying or exploring the behavioral attributes that might make them possible.

abundant evidence of positive rates of cooperation. Similar “anomalies” are found in voluntary contribution to public goods games (which are multi person Prisoner’s Dilemmas), where one sees positive contribution levels, in the trust game, where one sees positive transfers in both directions, and in many other instances.³

Game theory has faced similar predictive failures in its treatment of behavior in the real world. As John von Neumann argued (citations follow), its canonical behavioral assumptions predicted devastating conflict between nuclear adversaries.⁴ This has not happened, and the nonoccurrence of the “most spectacular event of the last half century” was the subject of Schelling’s Nobel lecture [2, p. 1]. Schelling could refer to this as an event—something which has taken place—even though it had not—because choice by self-regarding players predicted it so unambiguously. The reality, I will argue, is that because of the disjuncture between human behavior and the self-regarding assumptions often used in formal game theory, the latter offers little guidance, normatively or predictively, in thinking about behavior or strategy in a world of potential conflict.⁵

Before considering in more detail Schelling’s evolving acknowledgements of the limitations of game theory in understanding deterrence, it is important to reflect on exactly why the theory is so barren in terms of its implications for policy or behavior. The main reason can be stated simply. So long as agents are self-regarding and there is some possibility of destroying an adversary’s offensive capability and/or its will to retaliate, von Neumann was right to characterize nuclear confrontation as a Prisoner’s Dilemma.⁶ And, because of the almost unimaginable destructive power of nuclear

³ See Kagel and Roth [3], Camerer [4] or Field, [5–8] for more discussion. In the trust game A can anonymously give B some or none of an initial stake, which is multiplied in value in the transfer. B then may, but is not obligated to return as much as she wants to A. If self-regarding players are rational, there are no transfers in either direction.

⁴ One can object that the unitary actor assumption is simply inappropriate when thinking about interactions among states, although one can also object that the approach is inappropriate when applied to individuals (see Thaler and Shefrin [9]). Two points are indisputable: first, von Neumann argued (and believed) that superpower confrontation was a PD, and second, if that was indeed the game, it did not end with the Nash equilibrium. Von Neumann was a pioneer in developing game theory as well as nuclear weapons, and this has resulted in a tension which can be resolved in one of two ways. The first is to argue that nuclear confrontation was not a PD, in other words, that von Neumann *did not know what he was talking about*. The second approach, adopted here, is to accept the PD metaphorically as representative of superpower confrontation, but to argue that the behavioral assumptions that drove von Neumann’s (and many other’s) thinking were flawed. The central premise of this paper is that the reason we did not and have not experienced nuclear annihilation is that evolutionary history has endowed most humans with predispositions against playing defect in a PD that might well end up being played only once (Field [5,8,10]). People (and states) do indeed sometimes defect. But even when the logic of a strictly dominant strategy is fully understood, *individuals frequently choose not to play it*.

⁵ Developers of formal theory have not been particularly concerned about this, placing more weight on logical consistency and theoretical novelty than on empirical validity.

⁶ A long tradition in the deterrence literature objects, and instead treats nuclear interaction as a game of chicken (see Zagare and Kilgour [11, p. 18]). Chicken involves A threatening to harm B in a way that will also damage A unless B backs off. The best response for either party is to back off in the face of a threat, but if both choose to escalate, the worst (least preferred) outcome ensues for both. Von Neumann did not see nuclear interaction as a game of chicken. Words were cheap. He did not argue that we should try and intimidate the Soviets by *threatening* to attack. He argued for attacking, and for attacking now. Those who reason in this manner tend to downplay or dismiss fears of retaliation,

weapons, particularly thermonuclear weapons, it is a PD that will be played only once if the Nash equilibrium is realized on the first iteration.⁷

In the Prisoner's Dilemma played once, defect (which in this instance means preventive war, preemption, or first strike) is the strictly dominant strategy for both players. As von Neumann argued, it is the only strategy a rational self-regarding player, assuming he is playing against a similar adversary, can choose.⁸ But it is evidently not the strategy chosen by either the United States or the Soviet Union through the four decades of the Cold War. For both sides, defection was trumped by a policy of restraint on first strike coupled with the threat of limited or massive retaliation,⁹ and this was true throughout both the atomic and thermonuclear eras and in spite of substantial shifts over time in the strategic balance between the two adversaries. How and why did this happen, and why did it prevent nuclear war?

From archival sources and interviews conducted by political scientists and historians we know a good deal about discussions that took place in the United States at the highest levels during the first two decades of the nuclear age, and we are learning more about similar debates that took place in the

since self-regarding agents would never retaliate *ex post* (as opposed to threatening to do so *ex ante*, which would not be credible). A large literature attempts to solve this problem essentially by assuming it away [11, ch. 2].

⁷ Preempters like von Neumann saw the nuclear standoff as a Prisoner's Dilemma in which the strategy space was limited to "attack immediately" or "wait." No truly self-regarding player, reasoned von Neumann, would ever wait. Doing so exposed the actor to avoidable risk and granted a benefit to the adversary (the continued option of preemptive strike). Wait turned the PD into a game of trust: the question then became whether the restraint would be reciprocated. Actual political actors, who are human, often do wait, and establish expectations of reciprocity through diplomacy or other means that sometimes are realized. Von Neumann had little interest in wasting time on the dynamics of deterrence, or why people might make transfers in trust games. His counsel was to launch now, and his argument is unassailable if parties are indeed entirely self-regarding. If one finds this conclusion unpalatable something has to give. This paper argues (and Schelling suggests the same), that deterrence worked and von Neumann's predictions failed because humans are not entirely self-regarding (a position anathema to those who consider themselves hard-headed realists). If that is so, behavioral science will be messier than economists and game theorists might prefer, because the ways in which human predispositions differ from the self-regarding baseline are not deducible from a simple set of first principles. Ideally, the deviations represent empirically validated generalizations from experimental or observational data. That is one of the premises underlying the growing field of behavioral economics, and accepting it means that research methods cannot be limited to blackboard economics.

⁸ Von Neumann assumed that a rational player was by definition self-regarding and he had little tolerance for those who might suggest otherwise. A more general framework for rational choice theory is this: individuals have stable and complete preferences, these preferences are transitive and independent of irrelevant alternatives, and individuals use all available information in choosing action that maximizes a utility function derivable from these preferences. To say that players are both rational and self-regarding is to restrict the range of allowable preferences, and in particular to require that people prefer more material goods to less, and life over death. Suicide bombers may be rational and self-interested, but they are not self-regarding. Neither are those who give all their money to charity because it gives them a warm glow. The more restrictions one places on preferences the more it becomes possible to test the theory against actual behavior. The most general version of rationality, the simple claim that people act in satisfaction of their desires, is unscientific. It is unscientific because it is not possible to validate empirically: there are no observational or experimental data that could disprove it.

⁹ The doctrines of Massive Retaliation and Mutual Assured Destruction differ principally in the proportionality of the response.

Soviet Union. In the U.S., the central disagreements were between those inclined toward preventive war/preemption/first strike and those recommending policies of deterrence or containment.¹⁰ Support for aggressive preemption was remarkably widespread. It was not limited to a “lunatic” fringe. To provide a compelling rationale for deterrence, one of the objectives of the work Schelling conducted in the 1950s, was to weigh in on one side of a policy debate whose resolution had enormous real world implications.¹¹

Much of what Schelling had to say was based on introspection, casual empiricism, and common sense. To most citizens confronted with the realities of the conflict between the U.S. and the U.S.S.R, a policy of engagement and non-aggressive deterrence seemed intuitively more reasonable than one of aggressive preemption.¹² But nuclear strategists weren’t necessarily like “everyone else”: they prided themselves on asking tough questions, pushing logic to its limits, and, if necessary, thinking about the unthinkable.¹³ The problem faced by Schelling in trying to bolster the case for deterrence with game theory was that such theory, when coupled with the common assumptions that agents are both rational and self-regarding, provides stronger support for preemption or preventive war.

No one understood this better, or articulated it more forcefully, than John von Neumann, coauthor of the book that helped launch the American intellectual romance with these methods [14]). *Games and Economic Behavior* didn’t discuss Prisoner’s Dilemmas, which hadn’t yet been formally characterized. But von Neumann followed the subsequent literature on non-zero sum games and the equilibrium concept for non-cooperative games developed by John Nash [15]. The Prisoner’s Dilemma is a classic venue for applying the Nash solution concept, yet the Nash equilibrium in the one shot PD has long troubled economists because it is so clearly inefficient.

In the PD, you have two plays: you can cooperate or defect. Against the play of cooperate, defect is the superior strategy, and against the play of defect, defect is the superior strategy. Both parties would be better off if they both cooperated, but the logic of strict dominance is unassailable. If the

¹⁰ These positions are not completely irreconcilable, because one can threaten a nuclear strike to pressure an adversary to do (or not do) something other than simply not launch a nuclear strike against one’s own territory. Indeed, the U.S. relied on such a threat to deter a Warsaw Pact conventional thrust into Western Europe, and some wanted to use the threat to force the Soviets to do other things, such as get out of East Germany, or abandon their atomic weapons. The most aggressive preemption—a surprise attack without prior threats or attempts at bargaining, is nevertheless hard to classify as deterrence. The difference between traditional and aggressive deterrence seems to be captured in the distinction between deterring from and forcing to, although to deter successfully may be to force another not to do something it wants to do—such as attack you.

¹¹ It would be a mistake, however, to think that Schelling, or any of the other strategists or defense intellectuals, adopted a totally consistent position. More often than not they were simply of two minds about a problem, or moved sequentially between positions as they struggled with conundrums that remain with us today. His writings, however, unlike those of Bernard Brodie or William Kaufman, had little direct influence on high level decision making during the Eisenhower administration.

¹² At the start of the Korean War, in July 1950, only 15 percent of the American public agreed that the United States should “declare war on Russia now.” In September 1954 a Gallup poll asked, “Some people say we should go to war against Russia now while we still have the advantage in atomic and hydrogen weapons. Do you agree or disagree?” Only 13 percent agreed. [12, p. 100].

¹³ The reference is to the title of a book published in 1962 by Herman Kahn [13].

players are rational and self-regarding, we will be hard pressed to explain why such a decision maker should play a strictly dominated strategy.

Von Neumann believed that inasmuch as the US-Soviet standoff was a Prisoner's Dilemma, and inasmuch as both actors were rational and self-regarding, the only defensible policy was immediate attack [12, p. 100]; [16]. Since there was some chance of destroying an adversary's offensive capability and/or will to retaliate by attacking, the best course of action was to launch now. Many others argued in a similar fashion. As the Joint Chiefs of Staff maintained in 1947, "Offense, recognized in the past as the best means of defense, in atomic warfare will be the only general means of defense" [17, p. 77].

One reason the Cold War remained cold was that proponents of preventive or preemptive nuclear war lost arguments in the late 1940s, throughout the 1950s, and again at the time of the Berlin Crisis in 1961 and the Cuban Missile Crisis in 1962.

By the end of the 1960s, most of the dilemmas of the nuclear age remained and if anything had intensified, but key analysts had begun to lose faith in the promise of game theory to illuminate them. The limitations of these methods in providing real guidance to problems of nuclear strategy had become obvious, and doubts could no longer be so easily papered over with optimistic claims that future theoretical progress would remedy these deficiencies [12, p. 261]. After the 1960s strategic studies were less likely to claim to be advancing game theory through the study of nuclear policy, or to be using such theory to untangle such operational challenges as target selection.

The loss of faith in these methods did not mean, of course, that the conundrums vanished.¹⁴ The fundamental policy divide between those inclined to preemption and those inclined toward deterrence remains with us to this day. The world situation has changed since the 1950s, and especially since the early 1990s, with the breakup of the Soviet Union making it no longer as easy to identify adversaries or at least their location. But the heirs to von Neumann's way of thinking continued periodically to occupy prominent positions in the executive branch of the United States Government.¹⁵

Advocates of preemption always come armed with strong rhetorical advantages. It is a tough minded policy that can appeal to considerations of both opportunism and prudence. To advocate a policy of preemption is to ground policy in the inexorable logic of strict dominance.¹⁶ Rational, self-regarding agents must play a strictly dominant strategy, or, by definition, they are not rational.

The strength of the case for first strike, surprisingly, is, for proponents, not much affected by the military balance between the two adversaries. If a nation is stronger, the argument goes, it must strike first to crush the will, damage command and control, and eliminate as much of the retaliatory capacity

¹⁴ Nor did it mean that the elaboration of such models in academic communities ceased.

¹⁵ Setting aside the novel problem of threats from non-state actors, the United States faced a replay of the arguments for preemption/first strike against the Soviet Union in the 1950s as it confronted the prospect of a nuclear armed Iran or North Korea, or an unfriendly but nuclear armed Pakistan. And simmering dissension continued after the breakup of the Soviet Union over whether it was prudent to consider nuclear armed Russia as our friend.

¹⁶ A strictly dominant strategy is a strategy that provides a superior payoff, irrespective of the strategy selected by one's counterpart. In the one shot Prisoner's Dilemma, defect is the superior strategy whether one's counterpart cooperates or defects.

of the adversary as it can. And if it is weaker, it must strike first, to benefit from the element of surprise, to use one's assets before they are lost and, by destroying some of one's adversary's offensive capability, mitigate the damage from the almost certain incoming salvo. Why almost certain? Because one's adversary will have made the same calculations, and, whether stronger or weaker, also have concluded that striking first is best. It is a mistake to believe that altering the military balance in either direction will necessarily weaken the calls for preemption from those who favor it.

To advocate a policy of nonaggressive deterrence, in contrast, is to premise policy on human behavior that can be neither recommended nor expected in a world assumed populated by rational self-regarding agents. Deterrence works when both parties play strategies that cannot be defended as strictly rational for self-regarding players, and make inferences that their adversary is not entirely rational or self-regarding. One cannot get beyond von Neumann's case for preemptive war other than by acknowledging this. Advocates of deterrence can, however, point to the principal defect of an aggressive policy of preemption, and it is not a trivial one. As a practical and moral matter, it leads in a conflict between nuclear adversaries more or less directly and more or less certainly to the deaths of hundreds of millions of people.¹⁷

Von Neumann understood that if one wanted to provide a prescriptive justification for first strike, assuming parties are self-regarding, game theory was very effective. Or, if one wanted descriptively to explain why the world had been destroyed in a nuclear conflagration, game theory worked well. Outside of the classroom and world of working papers, in other words, theory could easily succeed at justifying policies that weren't pursued and explaining events that didn't happen. If one wants, in contrast, to justify a policy of deterrence and containment, explain descriptively why a balance of terror kept the peace for forty years during the Cold War, or provide guidance as to what one should do if one does not strike first, formal theory premised on rational choice by self-regarding agents turned out to be of relatively little help.

In the academic and policy worlds, however, those who could claim that they were using game theory to understand real world problems and to provide guidance on how to resolve them enjoyed a premium in the form of career advancement and honorifics. This is a sensitive issue but I think most will agree that it is a fact of intellectual politics that has been true for half a century and remains so today. What then was Schelling's attitude toward formal or pure game theory?

¹⁷ This, however, has never been a compelling argument for committed advocates of preemption, since victory is defined as retaining a higher fraction of the surviving population, territory, or economic assets. Obviously, if a country were able to obtain what it wanted from an adversary merely by threatening the use of nuclear weapons, it could be to that country's advantage to do so. But in order for threats to be credible, one must actually be prepared to follow through on them. In the age of conventional war, the victor could often be better off in spite of the costs of fighting. The problem for those pushing the conventionalization of nuclear weaponry, and the idea that one could fight and win a nuclear war, was that by most reasonable standards, this cannot be the case in a nuclear war. As Eisenhower put it, "even assuming that we could emerge from a global war as the acknowledged victor, there would be a destruction in the country (such) that there would be no possibility of our exercising a representative form of government for at least two decades at the minimum" (cited in Jervis [18, p. 62]).

Circa 1960, it can best be described as inconsistent. In the preface to *The Strategy of Conflict* he suggested that he was advancing game theory but later in the book gave mixed signals as to whether he believed such theory, at least at its then current stage of development, could provide practical guidance in matters of nuclear strategy or behavior. The Nobel citation in 2005 nevertheless awarded the prize to Robert Aumann and Schelling for contributions that the committee believed each of them had made to both theory and applications:

The work of two researchers, Robert J. Aumann and Thomas C. Schelling, was essential in developing non-cooperative game theory further and bringing it to bear on major questions in the social sciences. Approaching the subject from different angles—Aumann from mathematics and Schelling from economics—they both perceived that the game-theoretic perspective had the potential to reshape the analysis of human interaction. Perhaps most importantly, Schelling showed that many familiar social interactions could be viewed as non-cooperative games that involve both common and conflicting interests, and Aumann demonstrated that long-run social interaction could be comprehensively analyzed using formal non-cooperative game theory.

...Eventually, and especially over the last twenty-five years, game theory has become a universally accepted tool and language in economics and in many areas of the other social sciences. Current economic analysis of conflict and cooperation builds almost uniformly on the foundations laid by Aumann and Schelling [19].

Much of this was, perhaps by necessity, an exaggeration. Yet it was not entirely accidental that the Nobel committee would suggest, in the press release announcing the prize, and the prize citation, that Schelling, in his 1960 book, “set forth his vision of game theory as a unifying framework for the social sciences.” It is an interpretation that Schelling invited. In the preface to *The Strategy of Conflict*, he advertised the work as “a mixture of “pure” and “applied” research” [20], strictly situating it within the *theory of games* [20, p. v, his italics] and in the text made repeated efforts to incorporate the apparatus of game theory as it was then developed, including extensive discussion and matrix presentations of two person games in normal form.

Schelling’s contributions to formal theory are not in the same category as those of John Nash or Reinhard Selten [21]. I don’t mean that they are necessarily less or more valuable, simply that they are not in the same category. As Anatol Rapoport, one of the reviewers of Schelling [20] put it,

Dr. Schelling’s book is ... not therefore (to) be judged as a contribution to game theory, as a game theorist understands it, but as a contribution to the problem of linking game-theoretical concepts with other concepts in order to make possible more determinate normative recommendations to the decision maker... [22, p. 434].

There is little in *The Strategy of Conflict* that can be considered an advance in formal theory, and in later decades, Schelling didn’t claim otherwise.¹⁸ If, objectively, Schelling did not advance theory, his work nevertheless succeeded in creating an impression in the minds of many non-technical readers—and apparently the Nobel committee—that he had. That he did not push theoretical frontiers

¹⁸ As he said retrospectively regarding his 1960 book, “I don’t think I had any noticeable influence on game theorists, but I did reach sociologists, political scientists, and some economists” [26].

should not necessarily be seen as a criticism, since, even after another half century of development, formal theory provides limited insight into the types of problems with which he was most concerned (Walt [23]).

By 1960 Schelling had read Luce and Raiffa [24] in great detail [20, p. vi] and knew that the analytic methods they described were of limited prescriptive or descriptive value in studying the problems with which he (Schelling) was concerned. His acknowledgement of these limitations is to be found partly in Chapter 1, where he allows that models based on the assumption of rational (which he implicitly assumes to mean self-regarding) behavior may be a “caricature” of actual behavior [20, p. 5], but more extensively in Chapter 6:

... some essential part of the study of mixed-motive games is necessarily empirical. This is not to say just that it is an empirical question how people do actually perform in mixed-motive games, especially games too complicated for intellectual mastery. It is a stronger statement: that the principles relevant to successful play, the strategic principles, the propositions of a normative theory, cannot be derived by purely analytical means from a priori considerations ... There is consequently no way that an analyst can reproduce the whole decision process either introspectively or by an axiomatic method. There is no way to build a model ... with the behavior and expectations of those decision units being derived by purely formal deduction ... It is an empirical question whether rational players, either jointly or individually, can actually do better than a purely formal game theory predicts and should consequently ignore the strategic principles produced by such a theory [20, pp. 163–164].

In this and nearby passages, Schelling presages discussions of cognitive modularity (see Barkow *et al.*, [25]), work on dual selves (e.g., Thaler and Shefrin, [9]), and, more generally, the subfield that has come to be known as behavioral economics.¹⁹ And he suggested explicitly that experimental work was strongly needed to mitigate the deficiencies of pure theory: “It does appear that game theory is badly underdeveloped from the experimental side” [20, p. 165]).

An important question is whether subsequent experimental research mitigates the deficiencies of formal theory or simply ends up casting them in a harsher light. Clearly, at the time Schelling wrote *The Strategy of Conflict*, he hoped for the former. He held out the promise of advancing theory so that it could be more useful, writing for example, that “in international strategy the promise of game theory is so far unfulfilled...” [20, p. 10]), suggesting that it would be or could be fulfilled.

Here in clear focus is the conflict between Schelling’s ambivalent aspiration to recognition as a theorist, and the acknowledgement that theory, uninformed by behavioral research, and largely unconcerned with empirical validation, could not provide the foundation for a science of deterrence. On the one hand the advertisement of the book as a work in both pure and applied research, on the other hand the Chapter 6 recognition of the absolute limits of introspection or axiomatic methods. On the one hand, the acknowledgements that the promise of game theory in this arena was unfulfilled, on the other hand the apparent optimism that it was *so far* unfulfilled.

If one accepts the logic of Schelling’s comments in Chapter 6, that one can’t reproduce the whole decision process introspectively or through an axiomatic method, that behavior and expectations

¹⁹ “But in the mixed motive game, two or more centers of consciousness are dependent on each other in an essential way” [20].

can't be derived by formal deduction alone, then any hope for advancing a science of deterrence must rest on data: experimental, observational, historical. A science of deterrence had (and has) to be behavioral, *i.e.*, it has to rest on observations of human behavior.

But if progress had to come from the empirical/behavioral side, then Schelling's efforts to push forward theory were investments likely to have low yields in terms of advancing a prescriptive or descriptive science. He acknowledged this (*i.e.*, in Chapter 6) although, as we have seen, not strongly enough or consistently enough to pose obstacles to the Nobel committee's suggestion that he had "develop(ed) non-cooperative theory further" and award a prize based not only on advances in theory but also on its applications to "familiar social interactions".

2. The Evolution of Nuclear Strategy in the 1950s

If we accept Rapoport's judgment that *The Strategy of Conflict* made little or no contribution to formal game theory, we can ask a related question: what influence did Schelling's work have on the practical design of nuclear strategy? His writings provided general intellectual support for deterrence as opposed to preemption, and ultimately to efforts at arms control. At the same time, his emphasis on the manipulation of risk and the prospect of threatening or fighting a limited nuclear war in pursuit of political or strategic objectives can be read as support for an aggressive version of deterrence—brinkmanship—and a potentially dangerous form of saber rattling.²⁰

That said, he does not in fact appear to have had a great deal of influence on the formation of military or strategic policy in the 1950s. Partly this is simply because his government service was mostly in Democratic administrations and that under Truman involved foreign international assistance, not nuclear policy. Even though Schelling is often included in a pantheon of defense intellectuals alongside individuals such as Bernard Brodie, William Kaufman, or Albert Wohlstetter [12, p. 3]), it is actually rather hard to identify his footprints in the history of strategic debates in the 1950s at RAND or elsewhere.²¹ To appreciate this we need to delve more deeply into the actual history of those debates.

Proponents of first strike occupied far more than the political fringe, and it is clear that both Presidents Truman and Eisenhower considered seriously the prospect of preventive war or a preemptive nuclear strike against the Soviet Union.²² Neither, however, was prepared to approve an

²⁰ There is little evidence that either the Soviets or the Americans ever took this counsel to heart: "Rather than being implacable, irrational, or manipulative, states appear to be cautious, flexible, and generally loath to take precipitous action during intense crises" (Zagare and Kilgour [11, p. 228; see also pp. 29–30]; and Zagare [27, p. 113]).

²¹ Citation counts demonstrate that Schelling has had a strong influence on academic thinking about strategic policy, but that is not necessarily the same as having an influence on policy. In this limited influence, Schelling was not alone. Rosenberg describes the impact of strategic thinkers in these terms: "Although such conceptual work was important in shaping public perceptions, and occasionally influenced the thinking of high policymakers or strategic planners, it generally had little relevance in the 1945–1960 period to the pragmatic concerns of operational planners" [28, p. 10].

²² For the evidence, see Trachtenberg ([12, pp. 100–152]).

unprovoked surprise nuclear attack against the USSR,²³ and neither was enthusiastic about first use, except, in the case of Eisenhower, in the instance of an actual or imminent conventional attack by Warsaw Pact forces on Western Europe. Each, however, was ready to use nuclear weapons in retaliation for an attack on U.S. territory. U.S. policy was clarified in the Dulles-Eisenhower doctrine (announced in January of 1954) of massive retaliation to either an attack on the United States or a *conventional* Soviet thrust into Western Europe (or possibly other provocations).

The situation in Europe and particularly Berlin created the most practical and immediate concern. The U.S. lacked the conventional forces to thwart a conventional advance in these areas. Air Force war plans, reflecting the doctrine of massive retaliation, anticipated, as a retaliatory response, hitting the Soviets with everything the U.S. had, concentrating on cities. Although, in terms of military planning, Eisenhower was unenthusiastic about preemption except in the case of an imminent conventional attack across Europe, some within the Air Force and the strategic community continued to argue for dispensing with the requirement of a provocation, or maintained that the very fact that the Soviets had nuclear weapons should be considered provocation enough.²⁴ The arrival on scene of the far more powerful thermonuclear weapons did not end support for preventive war, which was widespread among civilian and military elites in the late 1940s and early 1950s.

In 1946 Leslie Groves, the U.S. army officer responsible for shepherding the development of the atomic bomb at Los Alamos during the Second World War, wrote an influential memorandum stating that

If we were ruthlessly realistic, we would not permit any foreign power with which we were not firmly allied, and in which we do not have absolute confidence, to make or possess nuclear weapons. If such a country started to make nuclear weapons we would destroy its capacity to make them before it had progressed far enough to threaten us [12, p. 100].

Groves' memorandum contained qualifiers about ruthlessness and reasonableness. Others dispensed with qualification, advancing proposals ranging from goal oriented saber rattling (threats) to massive surprise attack. In 1948, William Laurence, the New York Times' science correspondent,

²³ This prospect was explicitly rejected in NSC-68, which defined U.S. strategic policy in the 1950s, even though the document deliberately exaggerated the Soviet threat (Rhodes [17, p. 106]). NSC-68 was approved by Truman in August of 1950 and declassified in 1977. Truman had of course given the go ahead for dropping two atomic bombs on Japan in 1945. But it was not until 1948 that he allowed the military to proceed with plans for further use of atomic weapons, and he made it clear he was making no commitment that he would use them again. NSC-30, as cited in Jervis [18, p. 24].

²⁴ As Trachtenberg writes: "In the late 1940s, and well into the 1950s, the basic idea that the United States should not just sit back and allow a hostile power like the Soviet Union to acquire a massive nuclear arsenal—that a much more "active" and "positive" policy had to be seriously considered—was surprisingly widespread" [12, p. 100]. As Rosenberg reports, as early as 1947 "the final report of the JCS Evaluation Board on the Bikini tests had recommended that Congress be requested to redefine "acts of aggression" to include "the readying of atomic weapons against us" [27, p. 17]. But, as the public opinion data indicate (see fn. 12), the American public was not enthusiastic about preventive war. As a general rule, humans seem much more prepared to initiate, justify, or approve of attack in the face of provocation than in its absence. Of course, if one can define the mere possession of offensive weapons as provocation, or the possibility that one might acquire such weapons as provocation, the distinction blurs.

recommended an ultimatum to the Soviets: shut down their atomic plants or the U.S. would launch an all-out nuclear war.²⁵ Winston Churchill favored threatening the Soviets: get out of East Germany or the Western powers would destroy their atomic facilities. Leo Szilard pressed for preventive war against the Soviets, simply to wipe out their atomic capability, and at RAND, John Williams was a forceful advocate for similar action. The conservative political thinker (and former Trotskyite) James Burnham was as well. Most remarkably, so was Bertrand Russell, in an address at the New Commonwealth School in London on 20 November 1948. In August of 1950, Secretary of the Navy Francis Mathews gave a speech arguing memorably that the U.S. should become the first “aggressor for peace” [12, p. 117]. Even George Kennan, though he did not endorse initiating nuclear hostilities, mused that a war the Soviet Union stumbled into, before they had a massive arsenal, might be the best solution for the U.S. [12, pp. 103–104].

In the spring of 1953, Eisenhower considered but ultimately rejected the recommendation of a high level study committee headed by retired Air Force General James Doolittle that the Soviet Union be given a two year ultimatum: come to terms or risk global nuclear war. An Air Force study in August 1953 made a similar argument. In May 1954 the Joint Chiefs of Staff Advanced Study Group recommended that the U.S. consider “deliberately precipitating war with the Soviet Union in the near future”. Eisenhower would not go along with that, either, and in the fall of 1954, approved an updated National Security Paper which stated (as had NSC-68) that “the United States and its allies must reject the concept of preventive war or acts intended to provoke war” [28, pp. 33–34]. What the declassification of archival information has underlined is not that the U.S. refrained from attacking the Soviet Union (we knew that), but rather how close the country came to doing otherwise, and how forcefully and persistently advocates of preemption pressed their case.

After Dulles’s announcement of the doctrine of massive retaliation in 1954, and after it was clear that Eisenhower would not (except in the case of a Warsaw Pact conventional offensive in Europe) approve a preemptive nuclear strike, advocacy of preemption or preventive war became somewhat less open, but enthusiasm for it remained very strong, particularly in the Air Force.²⁶ The option of preemptive war was raised again by three members of the Gaither Committee in the fall of 1957 [28, p. 47]; see also [17] (pp. 106–108). Nor did the consideration of attack die with transition to new political leadership. In the 1960s, Kennedy seriously contemplated a nuclear first strike against the Soviets at the time of the Berlin crisis in 1961, and he was strongly pressured to resort to nuclear weapons in the Cuban Missile Crisis in 1962. From a military standpoint, the training

²⁵ Laurence was no ordinary journalist. Invited by Groves, he was the only reporter to witness the Trinity blast in New Mexico as well as the atomic bombing of Japan—he interviewed the pilots who flew the aircraft that dropped the bomb on Hiroshima and Laurence himself flew in an observation plane to witness the bombing of Nagasaki.

²⁶ Eisenhower was prepared to launch on warning either of a nuclear strike on the U.S. or a conventional attack on Europe. This can be viewed as preemptive, because attack would be initiated before bombs actually hit the US or Warsaw Pact tanks moved west, but not in the way advocates of preventive war meant it. The distinctions lie in whether the trigger was an immediately impending attack on the U.S. or key allies, the simple possession of nuclear weapons, or the mere possibility that they might be acquired. Prior to the introduction of ICBMs, anticipated warning times for an impending attack were days rather than minutes.

requirements for a strike in the presence of a conventional Warsaw Pact build up or in the absence of such direct provocation were essentially the same.²⁷

Nuclear weapons offered the prospect of a relatively cheap alternative to the conventional forces that would otherwise be required to repel a Warsaw Pact offensive westward across the European plain. The difficulty with this strategy was that the Soviets were capable (or at least the U.S. believed they were capable) of promising retaliation in kind. Because the prospect of flattened American cities as a result of retaliation might give American planners pause in the event of a Soviet offensive, U.S. defense intellectuals such as Bernard Brodie and William Kaufman argued early on that the threat to defend Europe in this way was not credible.²⁸

What were the alternatives? Giving NATO countries access to their own nuclear bombs was initially unpalatable to U.S. strategists, so critics of massive retaliation developed a competing doctrine—counterforce—emphasizing a different targeting strategy and more graduated escalation.²⁹ Counterforce entailed responding to a conventional incursion into Europe with a *limited* nuclear response against Russia: hitting airbases, silos, and military installations, but not cities. Perhaps the Soviets would respond against similar targets in the U.S. but, it was argued, they would see their interest in avoiding our cities since we had avoided theirs. We might then threaten to take out their cities one by one until the war could be concluded without it having escalated to an all-out nuclear exchange.³⁰

As an alternative to massive retaliation, the counterforce strategy met with initial resistance from the Air Force, where it was seen, essentially, as soft on Communism (and Communists). In June 1958 Air Force chief Thomas White told an audience of national security specialists he was “disturbed” by the recent tendency “to consider seriously self-restraints in nuclear weapons planning in the face of sure knowledge that no such restraints will be applied by the enemy. Our preoccupation with niceties in nuclear warfare... would, I am sure, delight the Kremlin.” Two years later, however, he supported the strategy. As far as we can tell, this was not the consequence of anything anybody at RAND had written. Why the change?

It had to do with the fact that the Navy’s Polaris program threatened the Air Force’s mission and budget. Submarine launched ballistic missiles (SLBMs), first deployed in 1960, had the great merit of being largely invulnerable, but they lacked the explosive power and accuracy to be used in a counterforce strategy. All they could do well was hit cities [30, p. 244]. Counterforce preserved a

²⁷ The main difference was how much lead time the military might have in preparing an attack, which would affect how many weapons could be fired off.

²⁸ This reasoning was part of de Gaulle’s rationale for pursuing an independent French nuclear deterrent, targeted at Soviet cities.

²⁹ Today both France and Britain have independent nuclear forces, and as McNamara argued in his 1962 Ann Arbor speeches, such forces make it even less possible to contemplate fighting a limited nuclear war (see Jervis, [18, p. 102]).

³⁰ The ideas of controlled escalation, and competition in risk taking, with the corollary that nuclear conflict might be limited, are significant features of Schelling’s thinking, although they conflict potentially with the viability of a nuclear firebreak, a principle endorsed in Appendix A of *The Strategy of Conflict* and reaffirmed in the Nobel lecture. Moreover, whatever “rules” of limited warfare the Americans wished to play by, the Soviets always ridiculed the notion of using nuclear weapons for bargaining [29, p. 144].

role for the more accurate bombers and land based ICBMs³¹ that the Air Force controlled, and was thus a way to marginalize the Navy.

The deterrent value of the Air Force's land based deterrent had been questioned since the early 1950s by Albert Wohlstetter and others because of its alleged vulnerability to preemptive attack by the Soviets, and there had been fierce debates about the merits of hardening the aircraft—putting them in underground hangars—*vs.* dispersing them or simply purchasing more of them to insure that a credible second strike force would survive. Curtis Lemay, head of the Strategic Air Command, brushed aside these concerns. He was confident that the then secret U2 overflights of the Soviet Union would, under any contingency, provide him sufficient warning to get his planes fueled, armed, and in the air. SLBMs promised to solve the vulnerability problem once and for all, but in doing so they threatened the Air Force mission and its budgets.

As a consequence of Polaris, the counterforce approach, which the Air Force but not the Navy would be able to undertake (because of the lower explosive power and poorer accuracy of the Polaris warheads) now had more support within the Air Force. But it was still a hard sell within the Strategic Air Command. When in the winter of 1961 William Kaufman briefed Power, who had succeeded Curtis Lemay as head of SAC, Power angrily responded: “Why do you want to restrain ourselves? ... The whole idea is to kill the bastards... Look, at the end of the war, if there are two Americans and one Russian, we win.” To which Kaufmann replied, “you'd better make sure they're a man and a woman.” Power then reportedly walked out [17, p. 67]; [30, p. 246].

This heated exchange illustrates why the prospect of millions, perhaps tens of millions or hundreds of millions of deaths was, for proponents of preemption, not a compelling objection to first strike. Victory was understood, and continued to be understood, not in absolute but in relative terms: as retaining a higher fraction of surviving population or military/economic assets [18, pp. 59-61].³²

Faced with resistance in SAC, and lack of enthusiasm from Eisenhower, who feared, given the dynamics of what he would later call the military industrial complex, that moving to flexible response was an invitation to launching aggressive war, Air Force chief White ultimately gave a weak endorsement to the counterforce strategy. Early in the Kennedy administration, however, these ideas came into ascendance, along with the idea that conventional forces should be built up to avoid having to choose between the unpalatable consequences of responding to an attack with massive retaliation and the equally unpalatable alternative of doing nothing. But counterforce, which on the face of it seemed highly attractive to those appalled by the likely consequence of either first strike or massive retaliation, brought with it its own set of issues.

³¹ Technical change during the 1950s was making land and air based weapons smaller and more accurate.

³² Carl Kaysen, in the 1961 first strike plan formulated at the height of the Berlin crisis, formalized this criterion somewhat less colloquially: “Accompanying these assumptions is the notion that prevailing in a general war means coming out relatively ahead of the enemy. As an example, if the US has lost 20% of its industrial capacity and 30% of its people, but the Sino-Soviet bloc has lost 40% of its industrial capacity and 60% of its people, then the US, somehow or other, has won the war” [31, p. 13].

3. Targeting Coordination

In the late 1940s the U.S. Navy actually condemned nuclear weapons as “immoral” [30, pp. 232-233]. This changed when the service, along with the Army, had the opportunity to begin acquiring its own tactical nuclear weapons. By the end of the 1950s, with the multiplication of nuclear armaments under the control of three of the four armed services (only the Marines didn’t have them), the obvious need for more coordinated targeting generated pressure for a single integrated operational plan (SIOP).³³ The SIOP for the fiscal year 1962 (effective 1 April 1961) that emerged from the Eisenhower administration and was inherited by President Kennedy and Secretary of Defense McNamara anticipated, as had earlier war plans, a massive preemptive nuclear strike against the Soviet Union, Eastern Europe, and China in the event of an actual or even *impending* Soviet conventional attack on Europe.³⁴

SIOP-62 represented the most aggressive posture advocates of first strike could obtain in the absence of a President willing to approve preemption (in cases other than an impending conventional attack across Europe). The overriding imperative in this plan was to provide a first strike capability in the event of conventional incursion in Europe, and to insure that a massive second strike capability would survive any attack by the Soviets on the U.S. and, in the event the U.S. was attacked, actually be used. Military planners saw plan rigidity as a merit, not a defect. SIOP-62 did provide a range of options in terms of how many missiles would be fired; this depended on the amount of advance warning or preparation time that the Air Force would be given.

For all of the options, the intent was to kill as many Communists as possible as quickly as possible. With a one hour warning, the retaliatory plan anticipated firing the 1459 U.S. nuclear weapons kept on alert (these ranged from 10 kilotons to 23 megatons, for a total of 2.164 gigatons); conservative estimates were that 175 million Russians and Chinese would die. If the President gave the go ahead for a full preemptive strike (28 h advance notice required) all 3,423 weapons with a total of 7.847 gigatons would be launched, and 285 million Russians and Chinese would die. These casualty estimates did not include deaths in Eastern Europe or victims of fallout around the world, and they reckoned damage only from blast, neglecting destruction resulting from heat, fire, and radiation. A more comprehensive estimate placed likely casualties closer to 1 billion [17, p. 88]. The option for a full preemptive strike was labeled Plan 1-A, giving insight into the priorities of some of the war planners ([28, p. 6]; [30, pp. 269–272]; [31]).³⁵ Total American megatonnage reached its peak in 1960, although total number of warheads peaked in 1966 [17, pp. 89, 95].

³³ Each Air Force Command controlled the detailed plans for the use of its weapons, as did the Army and the Navy with its tactical weapons. With the Navy acquiring strategic weapons (Polaris), the situation was worsening, with little coordination and much duplication of targeting [30].

³⁴ In an earlier briefing on SIOP-62, before Kennedy took over, David Shoup, Commandant of the Marine Corps, asked SAC chief Power what would happen if the Chinese were not involved in the fighting. “Do we have any option so that we don’t have to hit China?” Power responded, “well, yeah, we *could* do that, but I hope nobody thinks of it because it would really screw up the plan” [30, p. 270].

³⁵ Kaplan had access to critical documents related to SIOP before they were reclassified under the Reagan administration (some have since been declassified a second time).

SIOP-62 mirrored the inflexibility of strategy which had been a feature of its predecessors, the Joint Outline Emergency War Plan, which evolved in the 1950s into the Joint Strategic Capabilities Plan, the Joint Mid-Range War Plan, and the Joint Strategic Objectives Plan, all of which were characterized by plans for massive retaliation or attack [28]. These plans had called for executing a massive retaliatory nuclear strike in the event of a “General War,” defined in the Joint Strategic Capabilities Plan as “an armed conflict in which Armed Forces of the U.S.S.R. and those of the United States are overtly and directly engaged.” The Army tried in 1958 to have the following words added: “as principal protagonists with the national survival of both deemed at issue” but the Air Force succeeded in having the amendment nixed [30, p. 277]. As a practical matter, the battle between preempters and deterrers was to define how sensitive would be the trip wire that would trigger massive retaliation. Provocation would still be required for the U.S. to launch, but the trigger would not have to be very substantial—a shooting incident in Berlin would have been enough—and the plans were so inflexible that they were in essence a Doomsday machine. Defenders—and one could find some support for this view in Schelling—argued that it was this inflexibility that made it such an effective deterrent.

McNamara and Kennedy initially wanted a wider range of options—the possibility of more flexible response. The successor operational plan, SIOP-63, reflected some degree of counterforce thinking. It divided Soviet targets into five separate categories, with the initial U.S. strike only on such strategic sites as air and missile bases and submarine pens. Under Eisenhower, the desideratum of inflexibility was reflected in the facts that Minutemen missiles had to be fired in groups of 50 or not at all and they were each rigidly preprogrammed to strike one target. The Air Force initially refused to reprogram its missiles to provide flexible targeting, acquiescing only after its funding had been cut for a month. McNamara’s deputy pointed out that if the military chiefs really believed in SIOP-62, there was little need for generals [30, pp. 280–281]. In some sense this was a rules *vs.* discretion battle; the generals worried that any flexibility weakened the deterrent power of nuclear weapons and possibly the resolve of civilian leaders to use them if necessary.³⁶

But as McNamara succeeded in replacing massive retaliation with somewhat more flexible response (although even the limited attack options in SIOP-63 involved massive megatonnage with huge civilian casualties from fallout), some of the defects of counterforce became more apparent. In August of 1960 the newly orbited KH-1 (Corona) satellites began generating photographs indicating that the missile gap, which Kennedy had successfully exploited in his Presidential campaign, like the previously touted bomber gap, was an illusion.³⁷ Eisenhower had been right in denying the existence

³⁶ The objection to flexibility is that it becomes an end in itself, a way of kicking the can down the road. With policy makers unable to decide in advance what they would do in various contingencies, the imperative became to preserve flexibility so that decision makers would be faced in the heat of crisis or battle with a choice among options they had not been able or willing to make with minds unpressured by immediate circumstance (see [18, p. 80]). The more flexible response is built into war plans, the greater the potential strain on communication and control, and the greater the likelihood that there will be no response, one reason the military tended to be averse to increased flexibility.

³⁷ The claim of a missile gap originated with the Gaither report in 1957, a panel set up by Eisenhower that included Paul Nitze, a Democrat, and as Richard Rhodes puts it, “went rogue” in terms of what Eisenhower had expected it to do. Nitze was in fact the main author of the report, according to McGeorge Bundy [17, pp. 106–108]. Rereading this

of such a gap, at least in favor of the Soviets. In fact, as the satellite reconnaissance data revealed, the U.S. had a 10 to 1 advantage in ICBMs. The Soviet strategic forces were so small and disorganized that counterforce now began to look appealing as a first strike option that might succeed in wiping out or very significantly degrading the Soviet retaliatory capability. The September 1961 National Intelligence Estimate concluded that the Soviets had just four operational ICBMs.

In 1961 Khrushchev was threatening a separate peace treaty with the East Germans and in August began to build the Berlin Wall. The American garrison in Berlin had enough food, fuel, and ammunition to survive without resupply for just 18 days. So, with very limited options if the Russians further tightened the screws in Berlin, Kennedy considered a preemptive nuclear attack on hardened Soviet missiles as an option. The plan, little known to this day, was developed by Carl Kaysen, who estimated that only five to thirteen million Americans might die [30, pp. 297–298]. The Soviet military capability would have been devastated, but they would still probably have been able to hit back at the United States with a few megatons, and “New York and Chicago, with their great concentrations of people, can be virtually wiped out by a small number of high yield weapons. In thermonuclear warfare,” Kaysen added, in case it was not apparent, “people are easy to kill” [31].

Counterforce, which had originally been proposed as a move away from the hair trigger (and possibly non-credible) policy of massive retaliation, now provided additional fuel for preempters as well as grounds for more nuclear weaponry: more powerful, more accurate, and more of them. If the Soviets were ahead of us, argued the Air Force, we needed more missiles and bombers to deal with our vulnerability. If they were behind us, as turned out to be the case, we needed more weapons to transform counterforce into a viable first strike capability and to stay ahead of them.

However—as a practical matter—as the U.S. continued to become strategically stronger, it could be argued that a counterforce targeting strategy increased the likelihood that the Soviets would launch a preemptive strike against the U.S.³⁸ Increased Soviet vulnerability, argued some, including Schelling, could paradoxically increase the threat to the U.S. If counterforce was destabilizing in this way, it would be better to go back to targeting cities, holding them, essentially as hostages.

McNamara, who initially endorsed counterforce because it offered the prospect of fighting a limited rather than an all-out nuclear war, now found that his success in championing it provided justification for even greater demands on the part of the Air Force for weaponry, which he wanted to restrain as he built up conventional forces. Although SIOP-63 was never altered to reflect this, McNamara soon cooled on counterforce, and began emphasizing the importance of Assured Destruction, eventually Mutually Assured Destruction. It was enough, he and Kennedy decided, if the U.S., having absorbed a Soviet first strike, could destroy a quarter of the Soviet population and half its industrial capacity. In 1962, Kennedy and McNamara decided unilaterally to limit US land

history is a reminder that the manipulation of evidence that Iraq possessed WMDs in 2003 did not represent the first time U.S. intelligence estimates and have been influenced by political and bureaucratic imperatives.

³⁸ This was particularly so in the early 1980s when President Reagan committed to a missile defense system. Although most are aware of how close the U.S. and U.S.S.R. came to nuclear war in the Cuban Missile Crisis, the two countries came extremely close again in 1983 because the Soviets, listening to the bellicose rhetoric of the Reagan administration, and observing the unprecedented U.S. peacetime military buildup, became convinced the U.S. was preparing for a first strike against them (see [17], [29, pp. 345–346]).

based ICBMs to 1054 (1000 Minutemen plus the existing 54 Titans), down from the 10,000 pressed for by SAC and the 3000 ultimately requested by the Air Force [17, p. 95].

By 1965 McNamara had come full circle, arriving at something close to the Eisenhower/Dulles doctrine of massive retaliation that had prevailed a decade earlier.³⁹ McNamara's disenchantment with counterforce and return to 1950s era strategic doctrine coincides roughly with what Trachtenberg [12] has identified as the beginning of the exhaustion of strategic thinking.⁴⁰

For the last half century we have, in a sense, been replaying old tapes. None of the issues involving how to make a threat credible, whether to target cities or military assets (counterforce or countervalue), whether we should prepare for assured destruction or limited conventional war, whether or not flexible response is desirable—none of these issues is new. All were identified and actively debated prior to the mid-1960s. The arguments raised in the 1970s and 1980s, ranging from those articulated by members of the Committee on the Present Danger (established in 1976) to the pressure to establish an “independent” Team B challenge to the CIA's estimates of Soviet capabilities and intentions (1976), to those that almost led to war with the Soviet Union in 1983, were the same as had been articulated earlier, with in many cases the same cast of characters or their protégés [17, pp. 124–126, 150–157]. It is symptomatic of the exhaustion of strategic thinking—and a tacit acknowledgment that game theory premised on rational self-regarding agents could not ultimately offer much help in formulating strategy—that after 1966, Schelling moved away from issues of nuclear policy (see, e.g., [32]).

Schelling's contribution to targeting debates during the heyday of strategic thinking is unclear, in part because he was often of two minds about many of the issues. He spent a year at RAND in 1959 and was thus connected to and familiar with the think tank out of which the ideas of graduated escalation and counterforce emerged. But Schelling was ultimately lukewarm toward these concepts. Certainly people like William Kaufman were more central in articulating and advancing the doctrine. Schelling ran war games at Camp David in September 1961, with Blue and Red Teams consisting of U.S. military strategists, including among others John McNaughton, Alain Enthoven, Carl Kaysen, McGeorge Bundy, and Henry Kissinger. The results, which might be considered an unusual form of

³⁹ US policy cycled back again to emphasize counterforce in the 1970s. Following Vietnam, Schelling moved away from direct involvement with the government and military policy. In 1970 he led a faculty delegation protesting Nixon's involvement in Cambodia; this action effectively ended his role as a defense intellectual. See Schelling [26].

⁴⁰ “Strategy as an intellectual discipline came alive in the United States in the 1950s. A very distinctive, influential and conceptually powerful body of thought emerged. But by 1966 or so, this intellectual tradition had more or less run its course [12, p. 261]... there was an intellectual vacuum in the whole national security area. The economists, and people heavily influenced by their style of thinking were for a variety of reasons drawn into this vacuum. What they had was something very general, a way of approaching issues, rather than anything that in itself suggested substantive answers that went right to the heart of the strategic problem. Looking back at this body of thought as a whole, it is clear that the publication of Schelling's *Arms and Influence* in 1966 marked something of a climax. After 1966 the field went into a period of decline: the well seemed to have run dry, the ideas were by and large no longer fresh or exciting...” [12, p. 44]. I would go beyond this, and, to extend the metaphor, suggest that there never was water in the well: The methods were as barren of useful insights in the 1950s as they were acknowledged to be by the mid-1960s.

behavioral research, revealed rather striking inhibitions against going nuclear, even in a small way [30, p. 302].

Subsequently Schelling directed an interdepartmental group within the National Security Council which “examined certain long-range aspects of political military planning.”⁴¹ The report, *A Study of the Management and Termination of War with the Soviet Union*, completed a week before Kennedy was assassinated, looked at how the U.S. and Soviet Union might bargain to bring a nuclear war to conclusion under a variety of different scenarios [33].

Schelling’s writings, which emphasized controlled escalation and limited shots across the bow to advance national objectives but also stressed trying to keep war from getting out of control, were more consonant with SIOP-63, with its flexible response, than SIOP-62, with its inflexible massive retaliation. But Schelling was ambivalent about counterforce: Trachtenberg describes Schelling in 1960 as “being pulled in both directions”, indicating that he ultimately came down in favor of a controlled counter-population (countervalue) strategy, which some of his colleagues found cruel or bizarre [12, pp. 35–38]. Certainly counterforce allowed for graduated escalation, which Schelling favored, whereas if massive retaliation were to be viewed as entailing bargaining it was going to be a pretty short conversation.

Counterforce, on the other hand, risked blurring the bright line between conventional and nuclear weapons, which as emphasized in his Nobel lecture and Appendix A of *The Strategy of Conflict*, he valued. And one could argue, consistent with Schelling’s writings, that the inflexibility of massive retaliation/cities only/mutual assured destruction made the policy a more effective deterrent and thus, arguably, contributed to preventing war. The more inflexible the response, however, the greater the danger from false alarms (false indications of incoming missiles or bombers). The merits of counterforce were nevertheless ambiguous and Schelling’s attitude towards it conflicted, as indeed McNamara’s came to be. Schelling’s support for a bright line between conventional and nuclear weapons was at odds with the elements of his thinking that emphasized bargaining, competition in taking dangerous risks, flexible response, and graduated escalation.

Schelling’s overarching framework, like that of Clausewitz, emphasized that war and diplomacy should be considered elements of a broad spectrum of bargaining behavior, and the emphasis on risk manipulation can be seen as support for a particular type of brinkmanship [12, p. 45].⁴² The scenarios worked through in the 1963 NSC report were illustrations of how this might work. Of course, neither Schelling nor anyone else has yet had actual experience fighting and bargaining within the context of a nuclear war. Alain Enthoven once shut down a general who questioned Enthoven’s expertise by noting that he had fought just as many nuclear wars as had the general [30, p. 254]. No one really knows how or whether one could bargain in a controlled way in the heat of threatening or actually exchanging salvos with a nuclear adversary. There is an enormous range of problems. For example, with civilian and military targets often only a few miles apart, how could

⁴¹ Under the Kennedy administration, Schelling participated in a number of interagency committees, including one that led to the establishment of the hotline between Moscow and Washington. See Schelling, [26].

⁴² There is some irony here given the general antipathy of Democratic policy advisors to what were sometimes seen as the dangerous and reckless policies of Dulles and Eisenhower.

one have counted on the Soviets to understand, in the face of a barrage of incoming missiles, that the attack was counterforce only, and thus be persuaded not to go after U.S. cities?

Although Schelling could explore the conduct of nuclear war only through simulations, he did have an opportunity to apply his insights to the waging of conventional war. The experience was not a happy one. He was asked in 1961 by Paul Nitze, then Assistant Secretary of Defense for International Security Affairs, to come to Washington as his arms control deputy. Schelling demurred, but recommended his friend John McNaughton, with whom he had worked in the early 1950s administering the Marshall Plan. McNaughton at first resisted, saying he knew little about arms and strategy, but Schelling promised to teach him everything he needed to know. McNaughton went to Washington, where he was instrumental in persuading the Pentagon not to block the limited Arms Control Treaty of 1963, which banned atmospheric testing. When Nitze became Secretary of the Navy in 1963, McNaughton became Assistant Secretary of Defense.

In 1964, he was charged by McGeorge Bundy with developing an “integrated political-military plan for action against North Vietnam.” The plan, based on Schelling’s ideas about how to wage limited war, was to use graduated escalation, employing large numbers of troops as a deterrent to the North’s invasion of the South and then applying pressure on the North through an air campaign. The idea was to wait for, or possibly invite a provocation from the North and then retaliate with a measured air campaign to force North Vietnam to change its behavior.

In planning that campaign, McNaughton visited Schelling to try and figure out the answers to several questions: what did the U.S. want North Vietnam to do or stop doing, how would bombing make them alter their behavior, how would the U.S. know this had happened, and what would prevent the North from reverting to what it had been doing previously once the bombing stopped? The two strategists were unable to come up with answers to any of these questions, although Schelling did advise McNaughton to limit the bombing to three weeks.

The campaign, code-named Rolling Thunder, failed to alter the behavior of the North Vietnamese, if anything, hardening their attitude, and solidifying their will [30, pp. 334–335].⁴³

4. Taking Data to the Theory

Thomas Hobbes’ work *Leviathan* (1651, [34]) has for decades been understood as the classic non-formal evocation of the Prisoner’s Dilemma. Schelling does not place the PD front and center in his 1960 book (there are a few references to it at the end) although it is central to the problems he explored. This centrality is recognized in the Nobel citation, which began with an evocation of the Hobbesian dilemma:

Wars and other conflicts are among the main sources of human misery. A minimum of cooperation is a prerequisite for a prosperous society. Life in an anarchic “state of nature”

⁴³ The problem with using threats or pressure effectively is that it requires that the entity threatened respond rationally. In contrast, successful deterrence requires a number of areas in which irrational logic and thought processes must prevail (refusing to attack in the first place; actually retaliating after deterrence has failed). The hubris of much strategic thinking comes in its confidence that one can know or specify in advance in what realms rational thought processes will and will not prevail, both in oneself and in one’s opponent.

with its struggle of every man against every man is, in Thomas Hobbes' (1651) famous phrase, "solitary, poor, nasty, brutish, and short".

The citation goes on to provide a capsule history of the development of game theory. But it is hard to see, realistically, how advances in pure theory have helped us, or are likely to help us in understanding the behavior of nuclear adversaries. One can read Schelling's position on game theory in *The Strategy of Conflict* as inconsistent, or characterize it as acknowledging its limitations but suggesting that with improvements it could do the job.⁴⁴ In his review of Schelling [20], Rapoport [22]) questioned the likelihood of this. He argued that the challenge was not simply to improve upon a slightly flawed approach:

These ideas, especially some of the striking paradoxes, are interesting and stimulating. I believe, however, that they indicate the necessity of transcending game theoretical thinking (i.e., thinking exclusively in strategic terms) rather than the need to incorporate into the theory of games matters which do not fit into its conceptual repertoire... The fact remains that there is no rationally justifiable conclusion that leads the two players of a Prisoner's Dilemma game without communication to insure for themselves the largest joint pay-off. Such an outcome can result only if "irrational" considerations are allowed to determine the choice of strategy, for example, "solidarity," "trust," "the determination to do the right thing, no matter what the consequences may be," etc. Such considerations have until now been anathema to the realists. Among the strategists, it is perfectly proper to advocate "calculated risks" based on bluff, blackmail, and intimidation, but risks based on trust (which admittedly may be misplaced, else the risks would not be risks) fall automatically outside the scope of strategy because the associated concepts are not even in the vocabulary of the strategist.

Rapoport is not always completely on target (it is now well understood, for instance, that the PD with communication is formally identical to the game without it) but his basic point is well taken. Schelling respected Rapoport—there is a footnote in *The Strategy of Conflict* referring to one of his earlier essays as magnificent [20, p. 7], but we may infer that this review struck too close to home. When Schelling reviewed Rapoport's *Strategy and Conscience* three years later, his tone was uncharacteristically harsh [35]. Rapoport's book is a sprawling affair, containing at its best some deep insights into probability theory, and some lucid critiques of the strategic way of thinking and its dangers, but one can legitimately criticize it from a number of angles. Nevertheless, one is struck in reading Schelling's review by the animus he appears to bear towards the author. Schelling appears quite angry, angry that Rapoport has tarred all strategists with the same brush, has failed to distinguish good theory from bad, and didn't have the courtesy to contact his targets to ask them what they in fact thought (rather than, perhaps defensibly, inferring it from their writings). My interpretation is that at the time Schelling felt compelled to defend those whose aim was to apply and improve game theory, a group in which he included himself. And yet, if one focuses on Chapter 6 of

⁴⁴ Still, if his position was that the theory lacked much predictive or explanatory power now, but might, with improvement, have it in the future, what was the justification for extensive inclusion of game theoretic apparatus in a book devoted to analyzing pressing current problems?

Schelling [20], Rapoport's and Schelling's understandings of and acknowledgements of the limitations of formal theory were not very far apart.

The referenced passages from *The Strategy of Conflict* and Schelling's review of Rapoport are the basis for the interpretation that Schelling was inconsistent in acknowledging or addressing the shortcomings of game theory as a prescriptive and descriptive guide to behavior, and conflicted about its prospects and his role as a theorist. These ambivalences, if anything, became deeper with time. In a symposium in his honor at Harvard University in October 2006 [36], Schelling observed that the Nobel citation gave him cause to reflect on whether he "was then or ever had been a game theorist." He recalled that he learned what theory he knows from Luce and Raiffa [24], a book he spent "more time studying than the Holy Bible," but noted that the Nobel committee seemed most impressed by work he had published before Luce and Raiffa appeared, and before he had read it.

This last is a very puzzling claim, since the Nobel citation states that Schelling received the prize for "developing non-cooperative game theory further", and it is hard to see how he could have made progress on this account prior to digesting Luce and Raiffa. Certainly, the avidity with which (by self-report) he consumed the work⁴⁵ suggests that he had high hopes it would be useful to him. By the end of 1950s, however, as evidenced in *The Strategy of Conflict*, he was aware of its limitations. That realization would have been unsettling to anyone schooled in the merits of rational choice analysis, and doubly troubling in a world in which rewards (honorifics, professional advancement) for apparently applying game theory to real world problems remained so high. In Schelling [20], any concern on this account was papered over by taking the optimistic position that yes, the theory was limited in its predictive or explanatory value, but future progress would resolve these deficiencies.

Formal game theory can be beautiful from an aesthetic perspective and challenging in terms of the mathematical and logical puzzles it presents. Social scientists benefit from a basic understanding of its concepts and principles if only to appreciate its limitations and because it has become part of the intellectual landscape. But if one is interested in solving real world problems, or using theory to understand human behavior, it is quite often a cul de sac. Ariel Rubinstein has probably been most forthright about this, arguing that formal theory simply should not be "a tool for predicting or describing real human behavior" [37, p. 616].⁴⁶ Rubinstein did not, however, necessarily conclude from this that research in game theory be shut down. A close reading is that his plea is that experimentalists and students of the real world stop harassing theorists with the disjuncture between theory and behavior.⁴⁷ Howard Raiffa (one of the coauthors of Schelling's "bible"), reached a somewhat similar conclusion, largely abandoning the theory program in his later career, and instead focusing on the practicalities of negotiation rather than its mathematical analysis (compare [24] with [38]). Anatol Rapoport said essentially the same thing as Rubinstein forty years earlier:

⁴⁵ He recalls that he spent between 100 and 200 hours reading it (Schelling, [26]).

⁴⁶ For evidence of similar attitudes in political science, in particular the assertion that logical consistency is more important than empirical validity, see Walt [39]).

⁴⁷ Economists often talk of taking theory (or a new model) to the data. The objection here seems to be to the efforts of behavioralists and experimentalists to take data to the theory.

The theory of games has been developed much beyond the zero-sum game, and it is not the fault of the theoreticians that the results are so frequently indeterminate or psychologically disturbing. The mathematical theory of games was never meant to be a behavioral theory, but only a mathematical one, which examines the internal logic of certain situations without necessarily drawing conclusions about what this internal logic may imply in human affairs [40, p. 437].

There can scarcely be any “problem in the social sciences” of greater import than that of understanding relations among nuclear adversaries. It is not surprising if what theory has actually delivered in this area has been a disappointment if the theory was never meant to be behavioral.

5. The Logic of First Strike

If agents are rational and self-regarding, there can be no logic of deterrence, only one of first strike. Von Neumann argued famously, with respect to the Soviets, “If you say why not bomb them tomorrow, I say why not today? If you say today at 5 o’clock, I say why not 1 o’clock?” ([12, p. 104]; Blair, 1957, cited in [16, p. 143]). As he neared death in 1957, he confided to his friend Hilary Putnam, a philosophy professor at Harvard, about his pessimistic view of the future. Von Neumann told Putnam he was “absolutely certain (1) that there would be a nuclear war and (2) that everyone would die from it” [41, p. 114].⁴⁸

Von Neumann’s analysis and conclusions were common at the time among intellectuals and those who had thought seriously about nuclear war, both before and after the arrival of thermonuclear weapons. Three years later, on the front page of the New York Times, the British Novelist C. P. Snow made a similar prediction, stating that absent massive disarmament, thermonuclear war within a decade was a “mathematical certainty.” As Schelling pointed out in his Nobel lecture, nobody at the time found this claim exaggerated [2]. Given what we know of the Emergency War Plan of the late 1950s or SIOP-62, or the Berlin crisis of 1961, or the Cuban Missile Crisis of 1962, or U.S.—Soviet tensions in 1983, this is not surprising. Von Neumann, had he still been alive in 1960, might have asked, why not sooner?

The failure accurately to predict has not inhered in the logic. The problem is with the implied behavioral assumptions. A taboo not codified in any treaty has kept nuclear weapons from being used against an adversary since 1945, just as a taboo kept poison gas from being used in the Second World War. In his Nobel lecture Schelling is rightly interested in these normative constraints on behavior and the degree to which they give rise to behavior at odds with what was predicted and counseled by von Neumann. Some behavioral inhibitions—in my view as much biological in origin as they are cultural—checked the pressures on or inclinations of leaders to attack. Our survival has depended upon the fact that decisions makers were and are, in some respects, the opposite of rational and self-regarding. It has depended in part on the fact that they were and are human.

⁴⁸ Von Neumann hated communists and had no moral qualms about working on the H Bomb project. From an early date he looked forward to nuclear conflict between the two superpowers. He expressed this anticipation in a 1951 letter to Lewis Strauss: “I think that the USA-USSR conflict will very probably lead to an armed ‘total’ collision, and that a maximum rate of armament is therefore imperative” [30, p. 63].

There is no practical defense against ICBMs. There is therefore, as Richard Rhodes puts it, “no military solution to safety in the nuclear age” [17, p. 101]. Game theoretic analysis, which commonly assumes that players are logical, rational, and self-regarding, leads to the conclusion that the surest and most effective way to reduce this vulnerability is to launch a surprise attack aimed at an adversary’s offensive weapons. Because players are human, however, and sometimes prone to retaliate to attack if they can, even if it gains them little or nothing, or even makes their deteriorated situation even worse, this is an almost certain invitation to mutual incineration. The effort to ground strategy in game theory produces, using Robert Jervis’ words, “doctrines which are incoherent and filled with contradictions, which though they are superficially alluring, upon close examination do not make sense” [18, p. 19].⁴⁹

In his Nobel autobiography, Schelling tried to interpret the convention against use of nuclear weapons as an example of a focal point. But this is not an appropriate use of the concept. Pure games of coordination have multiple equilibria, and the concept of a focal point refers to the role of tacit or implicit knowledge in enabling players to coordinate on one of them. The Prisoner’s Dilemma is not a pure game of coordination. It has only one equilibrium (unlike the New York meeting problem that Schelling popularized), and that equilibrium is inefficient, unlike any of the equilibria in a typical coordination game. The no-nukes “convention” is not a Nash equilibrium, as disgruntled advocates of preemption have pointed out again and again, and there are strong incentives, pressures, and temptations for players to deviate from it, in contrast to what is true for a coordination game equilibrium once reached. If there are norms (whether biological or cultural in their origin) which incline us to the cooperative solution in PDs, they are different from the social or cultural norms that contribute to a focal point. The former require in some respects that we be the opposite of self-regarding, whereas the latter (such as those, within a common language area, that allow us to agree on the meaning of a word) do not (Field, [42]). We are best served by acknowledging these differences rather than suggesting that these types of norms are of the same genus.

The analysis of games of coordination does help us understand the role of tacit or implicit knowledge in solving such problems as which side of the moving walkway to stand on, or where, in a pre-cell phone age, and in the absence of prearrangement, we should meet in New York. In that sense, Schelling has contributed to the use of game theory to understand “everyday economics”, to use Warsh’s words, or “familiar everyday problems”, to use the words of the Nobel citation.

But the economic and social significance of coordination problems pales in comparison to those presented by Hobbesian dilemmas, and the Nobel citation’s emphasis on Schelling’s contribution to the analysis of “non-cooperative games that involve both common and conflicting interests” makes it clear that they had in mind the latter category of social dilemmas. The reality is first that Schelling made little contribution to the formal development of non-cooperative game theory and second that

⁴⁹ He goes on to argue, “To argue that any nuclear doctrine must be at least partly irrational does not mean that all doctrines are equally at odds with the reality they try to reflect and shape. If one starts with misleading conceptions, the more complete and thorough the reasoning, the stranger and more confusing the results. Only by understanding and accepting the implications of nuclear weapons can we develop a more appropriate policy. But even such a policy cannot meet all the standards we normally require of rationality” [18, p. 20].

such theory is of little value in understanding how humans solve Hobbesian, or Prisoner's Dilemmas. On both of these counts, then, the Nobel citation reflects a certain amount of misdirection.

In the long run it matters little why someone wins the Nobel Prize. What does matter in the social sciences is correctly identifying what types of contribution advance our understanding of human behavior. The challenge of understanding what contributes to collective action or collective inaction (restraints on harm) beyond the family are central to social and behavioral science. Claims to advance must be carefully vetted, with the most important criterion not formal axiomatic consistency but the extent to which actual human behavior is explained or predicted. Misdirection with respect to the contributions of game theory has been facilitated by a finding of common ground among those who love the intellectual challenges it poses, even as they acknowledge its limited applicability to actual human behavior (e.g., Rubinstein), and a broader audience that understands enough about the approach to appreciate its apparent potential, but not enough to recognize its limitations.

Deterrence can only work if in some respect humans are not rational in their logic and/or self-regarding in their preferences. A strategy of defensive deterrence has two main pillars: (1) each counterparty refrains from first strike; and (2) will in fact retaliate in the event deterrence fails. Neither pillar, von Neumann appreciated, could be defended as the behavior of a self-regarding rational actor. To von Neumann, failure to strike first was irrational and foolish. It was foolish because it exposed one's country to the risk of being hit first. And since the nuclear counterparty had access to similar logic, and assuming the adversary was also a self-regarding rational actor, the attack must be imminent.⁵⁰ If missiles were not already inbound it was only because the adversary was also behaving in a foolish manner. To refrain from attacking first was to put one's faith in irrational behavior on the part of one's counterpart. And, if one's counterpart were behaving in a foolish fashion, failure to strike first meant passing up an opportunity to gain at her expense. In a thermonuclear world, to refrain from first strike was to refrain from defecting in a Prisoner's Dilemma game which would have more than one play only if both parties irrationally chose not to defect.

Some have explained the failure of the United States to wage preventive war in the early 1950s as due to its self-perception as relatively weak [12, pp. 100–115]. Yet in 1961, when intelligence revealed the disarray and weakness of Soviet strategic forces, and Khrushchev's nuclear bluster had been shown to be largely a bluff, the U.S. again passed on the preemption option. Conditions were more favorable for a U.S. first strike than they had been at any time since the late 1940s. And in his September 1961 war games, Schelling could not get either the Blue or the Red teams to go nuclear, no matter how hard he provoked. In the policy domain, the 1961 Kaysen plan at the height of the Berlin crisis was rejected, even as a contingency [30, pp. 299–300]. Something obviously restrained the willingness of even tough minded strategists like Kaysen and Kissinger to act on the logic of first strike. Von Neumann must have turned over in his grave.

As many have pointed out, the second pillar of deterrence, the promise of retaliation, also requires behavior that cannot be justified as rational. Having decided (irrationally) to refrain from first strike, the threat of retaliation is intended to deter a counterparty from doing the rational thing, which is to

⁵⁰ See Kahn, ([13], pp. 151–152, cited in Rapoport, [40], p. 134). Kahn describes an imagined conversation between a Soviet general and Khrushchev as they consider the pros and cons of a first strike on the United States.

attack first. But if the attack has already happened, deterrence has failed. With many U.S. cities in ruins, of what possible value would have been the destruction of millions of Soviet lives, simply to prove, after the fact, that we had “meant” what we said when we promised/threatened retaliation?

The threat of massive retaliation would not be credible if the counterparty believed one were completely rational because it would make no sense to carry through on it after the fact. This train of logic leads to the conclusion that a rational actor should not and would not be deterred from attacking by a threat of retaliation so long as she assumed the target was rational, *because a rational victim of aggression would not retaliate*. It is easy to threaten retaliation, but talk is cheap. This reality helps illuminate the obsession of military planners with inflexibility, as they have struggled over the years to make retaliatory threats credible.

Schelling argued that madmen could not easily be deterred by threats of retaliation [20, p. 6]. An apparent corollary is that those who are rational (not mad) *can* be deterred. But as this analysis reveals, you cannot deter a rational actor who also believes you are rational, because your threats of retaliation will not be credible.

This issue exposes one of the soft underbellies of strategic thinking, and is one that Schelling chose not to probe too deeply. In *Arms and Influence* [43] he cited Max Lerner, writing in *The Age of Overkill* (1962) that “The operation of the deterrence principle in preventing war depends upon an almost flawless rationality on both sides.” ([44, p. 27], cited in Schelling, [43], p. 229). Without directly challenging or contradicting Lerner’s claim Schelling dryly notes that “...when people say that “irrationality” spoils deterrence they mean—or ought to mean—only particular brands of it” [43, p. 229].

The practical success of deterrence rests in part on the fact that the very act of refraining from first strike, which is demonstrably *not rational*, may increase the credibility of a threat of retaliation. By contravening the counsel of von Neumann and other aggressive preempters, an actor illustrates by her restrained behavior a willingness to behave in an irrational although still apparently instrumental fashion. Thus a counterparty might, for this reason, think twice about dismissing the threat of massive retaliation on the grounds that such retaliation would, after the fact, be irrational.

Strategists like von Neumann argued that a nuclear war initiated by a U.S. first strike could be won, if the U.S. possessed sufficient strategic superiority, and this argument surfaced again and again in the 1960s, 1970s, and 1980s. There were several targeting strategies to choose from. One could focus principally on military targets, to degrade or eliminate second strike capability. Or one could focus on CCC—command, control and communication—decapitation strikes that killed leaders and destroyed communications capability, so that even if weapons survived and decision makers were (irrationally) angry and prepared to retaliate, the weapons would not be fired. Or one could stick with the traditional cities only strategy, demoralizing the Soviets (“Shock and Awe”) to make sure that there was little or no (irrational) will or desire left for retaliation. Advocates of preemption stressed that in order for the U.S. to face a risk of retaliation, both the Soviet capability and the will to retaliate had to survive our first strike. And even if there was some probability that both did survive, winning the war could be defined as the U.S. suffering lower relative losses.

Success for the U.S. as aggressor would require combining an assumption of rational first strike on our part with irrational restraint on first strike on the part of the Soviets (otherwise the Soviets

would already have attacked and U.S. would not have been able to strike first) followed by rational behavior by the Soviets following our strike (why bother retaliating since their effort to deter us had obviously failed). If, in contrast, we believed the Soviets would be consistently irrational in their behavior (irrational restraint combined with irrational retaliation in the event we attacked) we might, if we were rational, be deterred from attacking.

But note that the non-aggression outcome associated with symmetric adoption of mutual assured destruction required asymmetric assumptions about the rationality of each of the adversaries. In particular, each had to believe that they were rational but their adversary was not—which is objectively impossible.⁵¹

To navigate in these clouded waters we will be less helped by game theory than by an understanding of social psychology, the theory of mind,⁵² and such regularities as the fundamental attribution error.⁵³ Contrary to Lerner's argument, successful deterrence requires not flawless rationality but elements of irrationality—sometimes real, sometimes perceived—on the part of both parties. It is likely that the predispositions that allow deterrence to work have in part a biological foundation (see Field, [5–8,10]).

If the scenarios for aggressive war pushed by von Neumann and others seemed to gamble with the nation's fate (what if the assumptions were off by a bit; was it OK if “only” five to thirteen million Americans died?), preempters had a ready answer. Their proposals were driven not just by opportunistic calculations, but by fears that the Soviets were making exactly the same calculations, planning to fight and win a limited nuclear war (see, e.g., Kaysen [31], or arguments advanced by the Committee on the Present Danger in the 1970s and 1980s). And if that were true, it was clearly in the U.S. interest to attack first. Would those who hesitated prefer confronting the aftermath of a Soviet first strike on the United States, with degraded hardware, disrupted command and communication, and/or devastated cities? From this it's a short distance to von Neumann's if you say 5 PM, I say why not 1 PM. The U.S. military was acutely aware of vulnerabilities in its own command and control systems, and resisted modifications in SIOP-62 on this account. It wanted inflexible response. The concern was that even if U.S. weapons survived a Soviet strike, the relevant individuals⁵⁴ might not be able or willing to launch them.

⁵¹ Robert Jervis was exploring similar issues when he wrote, in 1984, that “A rational strategy for employing nuclear weapons is a contradiction in terms” [18, p. 19].

⁵² Theory of mind involves such questions as how animals (including humans) infer another's knowledge or intentions from the direction of its gaze, the timbre of its voice, or the expression on its face (for discussion, see Cheney and Seyfarth [45]). In game theory, none of this matters: there is promise of a parsimonious short cut to inferring intention. Sometimes, but by no means always, this short cut provides good predictions, but where it fails, and when the setting is nuclear confrontation, the failures are particularly problematic.

⁵³ The error, which has been widely demonstrated, is to assume that my behavior is governed by the situation while yours is governed by your disposition, in short, to assume that I'm rational and you're not. Since the success of mutual assured destruction requires mutually asymmetric assumptions about the rationality of the two parties, deterrence could not succeed in the absence of this cognitive bias. The classic study is Ross [45]).

⁵⁴ It would be perhaps a misnomer to call them decision makers, because the whole point of inflexible response was that they were not supposed to be making decisions.

Schelling, of course, preferred to win without fighting, or by fighting as little as possible (“a theory of deterrence would be in effect a theory of the skillful nonuse of military forces” [20, p. 9]). But he acknowledged at least implicitly the asymmetric but complementary roles rationality and irrationality would have to play in successful deterrence, although they were the reverse of what was needed for successful first strike.

Schelling repeatedly emphasized the potential role played by irrationality:

“It is not a universal advantage in situations of conflict to be inalienably and manifestly rational in decision and motivation. Many of the attributes of rationality, as in several illustrations mentioned earlier, are strategic disabilities in certain conflict situations” [20, p. 18].

But he does not confront the theoretical incoherence of a model that requires that each adversary believe that they are rational whereas their adversary is not, or that they are rational but must make their adversary believe the opposite, or that they are or will be rational in some instances but not others. We have already seen that a rational attacker (A) would not find credible the threat of retaliation by B if A attributed rationality and self-regarding preferences to B. A cannot be deterred by B if she attributes rationality and self-regarding preferences to B. If A is rational and self-regarding and has not attacked, it must be because of her attribution of some degree of irrationality to B (in other words, A must have some fear of retaliation). But if that is so, the merely *threatened* use of nuclear arms by A in an attempt to alter B’s behavior in ways distinct from simply refraining from attacking would also lack credibility. If A is deterred today by B’s irrational threat of retaliation, why should that be any different tomorrow? Aggressive preemption or the threat of aggressive preemption by A will succeed only if B combines irrational restraint prior to A’s attack (else why would A have the opportunity to strike first?) with rational calculation of B’s self-interest following A’s strike.

If we are wedded to a rational choice modeling approach, should we not expect at least a consistency in the irrationality of B before and after A’s rational strike? Experimental research makes it abundantly clear that humans are inclined to have their behavior influenced by both rational and irrational thought processes, and it sometimes makes a big difference which prevail. *Formal game theory offers us no help in estimating those probabilities.* If your fate and the fate of the world depended upon it, would you rather your political leaders be good intuitive psychologists, skilled at inferring the emotions, motives, beliefs, and intentions of counterparties, and of themselves, or well-trained game theorists?

To accept the argument that nonaggressive deterrence (peaceful coexistence) requires a combination of irrational and rational behavior on the part of both parties is not to reject such strategies as worth pursuing, nor is it to argue against being deterred in the face of a determined and committed opponent. Such policies and behavior often are worth advocating and pursuing, and as an empirical matter often have worked as they were intended. The problem for game theory is that in a world of self-regarding agents, there is always a stronger (more hard-headed, more tough-minded) case to be made for preemption. Allowing oneself to become too absorbed in the strategic way of thinking runs the risk of finding oneself in a race to the bottom, in which aggressive policies systematically trump those that are less so.

Leslie Groves' advocacy of a policy of destroying the relevant facilities in any country that might threaten to develop nuclear weapons sounded much like Vice President Cheney's advocacy of a broad license for preemption more than half a century later: "If there's a 1 percent chance that Pakistani scientists are helping al-Qaeda build or develop a nuclear weapon, we have to treat it as a certainty in terms of our response. It's not about our analysis, or finding a preponderance of the evidence. ...It's about our response" [47, p. 62]. Since a 1 percent probability is low, its estimation subjective, and one Cheney proposed the United States arrogate to itself, there was little limit, if this doctrine were accepted, to the range of targets an attack upon which could still be clothed in the rhetoric of defense.

What is troubling about a policy of preemptive or preventive war, as applied to Iraq, and possibly to Iran or North Korea, are the implications for the U.S. if it is adopted by other nations. What it means is that any nation, on its own say so, would be justified in attacking the United States or any other country if it judged that there was even a very small probability that the United States might attack it. From a standpoint of rational choice theory and a model of a world consisting of self-regarding agents, there is no error in this reasoning.⁵⁵ International law has always recognized the right of a nation to protect itself by attacking an adversary when an attack from that adversary is "imminent".

The Cold War spanned four decades, and its epicenter can be said to have been Berlin. The conflict began in 1947, two years before the blockade and airlift. Conflict over the city brought us close to nuclear war in 1961. The Cold War ended with the fall of the Wall in 1989 and the subsequent collapse of the Soviet Union (1991). For U.S. baby boomers the Cold War was an omnipresent fact of life, a feature of it from their birth until, to the surprise of many and consternation of some, the conflict suddenly ended a quarter century ago. The Cold War defined the political, strategic, and intellectual milieu within which Schelling's most influential work was conducted.

Over most of that period, a philosophy and strategy of containment and deterrence helped prevent thermonuclear war between two well-armed adversaries, the Soviet Union and the United States. Neither launched a surprise attack on the other. The question I have asked in this essay is whether the success of the Cold War stalemate depended upon the rational and self-regarding behavior of both parties, and on each party's belief that the other was rational and self-regarding. If it did, then the interaction of these two adversaries would indeed be a fitting subject for game theory. Lerner claimed that it did and it was. Schelling danced around the question. To challenge the assumption made by Lerner that successful deterrence required consistent (flawless) rationality on the part of all parties too directly might have undercut the impression that game theory/rational choice approaches were central to the conclusions of his analysis.

That said, Schelling appears to have recognized the flaws in Lerner's claim. Without acknowledging the poor predictive power of models premised on rational choice by self-regarding agents, in

⁵⁵ Note that a similar issue applies to the abandonment of the Geneva conventions, which former Attorney General Alberto Gonzalez characterized as "quaint." There is little reason from a game theoretic perspective why the Geneva conventions should be respected, but clearly to some degree they are and have been, and many in the U.S. military objected to the US adoption of methods of torture, which resulted in the deaths of tens of captives, on the grounds that such actions threatened the protections available to captured American soldiers.

particular the fact that individuals often do refrain from first strike, even when rational analysis indicated, as von Neumann protested repeatedly, that it is in their interest to do otherwise, and that they often will respond in kind to attack, even when this makes no sense *ex post*, there would have been no explanatory space for the topics that interested Schelling. There would be no arena for the murky world of threats and promises, driven by conflicting behavioral predispositions within individual humans themselves and characterized by the asymmetrical ability to benefit others by failing to do them harm as opposed to providing them with affirmative assistance. It is his acknowledgment of that world of shadows, in which things were not always what they seem, and what seemed to be the pursuit of flawless rationality was sometimes nothing of the sort that makes Schelling's work interesting.⁵⁶

6. Rational Choice

The Royal Swedish Academy of Sciences saw it differently. Rather than interpreting Schelling as having shown that behavior interpreted as the pursuit of flawless rationality was often nothing of the sort, the Academy's prize committee claimed that Schelling (and Aumann's) contribution was to show the opposite: "A consequence of these endeavors is that the concept of rationality now has a wider interpretation; behavior which used to be classified as irrational has become understandable and rational" [19, p. 3]. This is misdirection. It reflects the same kind of papering over of contradictions which renders most discussions of U.S. strategic policy incoherent. Human behavior is what it is, and sophistry cannot make the play of a strictly dominated strategy rational. If *x* is black, and we understand what black means, one can't make it white simply by reinterpreting it as white.

In describing behavior, the word rational is often used in different ways. The broadest meaning is simply the claim that people act in satisfaction of their own desires. This version is not interesting from a scientific standpoint, since it is impossible to conceive of any behavioral data that could not be made consistent with it. A narrower version posits that people have goals reflected in preferences, that these preferences are both stable and transitive (if *A* is preferred to *B*, and *B* to *C*, then *A* is preferred to *C*), and that people use all available information to choose a course of action likely to lead to the realization of these goals.

The third and most rhetorically powerful version of rational choice carries forward all of the language about goal seeking, but adds that preferences are, in addition to being stable and transitive, also self-regarding. Following Gintis [48], I suggest that self-regarding is superior terminology to

⁵⁶ Schelling's work is as much about diplomacy as the art of war fighting (as witness the title of his 1966 book: *Arms and Influence*). But, as those skeptical of diplomacy have always argued, talk is cheap, and actions speak louder than words. If these precepts are taken to heart, diplomacy is not part of the strategy space. As von Neumann understood, so long as players are rational and self-regarding, the strategy space was limited to launch or not launch, and the case for the former was unassailable. The idea that threats or promises might be used to further political or strategic aims presupposed that actors would behave in ways that could not be defended as rational. A country might threaten to retaliate, or threaten to build a Doomsday machine to overcome the prospect of weakness of will, but a rational actor would never do either. The idea, popularized by Schelling, that, short of attacking, one might create a situation where things might get out of control and an attack "might happen" would have struck von Neumann as mealy mouthed. "Just do it" would have been his reaction.

self-interested, since it avoids the possible ambiguity arising when individuals experience a “warm glow” from helping others. Such helping behavior (and I do not mean to suggest that all helping behavior has this characteristic) might arguably be self-interested, but is not what we would mean by self-regarding.

In using the term rational in this paper, I have meant this third, most restrictive, and most rhetorically powerful use of the term. Under this meaning, to assume rational choice is to assume that people (or countries) act so as efficiently to advance their material self-interest.

7. The Logic of First Strike—Once More with Feeling

It is a commonplace among tough minded thinkers to say that talk is cheap, and actions speak louder than words. But rhetoric matters. It affects human attitudes and behavior, even if theory often suggests it shouldn't. Most humans are inclined to support military operations in response to what is seen as an unprovoked attack. Public opinion tends to be less comfortable with preemptive attack, because the threat countered is probabilistic and thus speculative. Aside from the small fraction of humans with sociopathic tendencies, attack simply for personal or territorial aggrandizements is the toughest sell of all: most humans are repelled by it. Even the most cynical dictators find it necessary and desirable to justify aggression in terms of a prior litany of perceived wrongs. For this reason Hitler felt compelled to cloak his invasion of Poland in 1939 as preemption to counter what he claimed was an intolerable threat from that country. The immediate *casus belli* was a ginned up raid on Germany supposedly conducted by Polish forces.

In 1949 there was not a 1 percent probability that the Soviets were obtaining the bomb: there was a 100 percent chance that they had it. But surprisingly, this is beside the point in terms of the argument for attacking now. The logic of first strike is inexorable for an analyst consistently applying “realist” assumptions. It does not depend much on the size of this probability. After all, one could imagine a hypothetical conversation between von Neumann and former Vice President Cheney, “If you say attack if there's a 1 percent probability, I say why not if there's a 0.1 percent probability. And if you say 0.1, I say why not 0.01.” Because of the asymmetrical ability to benefit others by not harming them, as opposed to providing affirmative assistance,⁵⁷ an asymmetry which is grossly magnified in the case of weapons of mass destruction, one requires very low probabilities in order to make the case for preventive war on these grounds.

Now let us again consider again the second pillar of MAD: the promise or threat of retaliation. In order for such a threat to play any role in deterrence, it must be credible. This is, of course, a persistent theme in Schelling's writings, and much subsequent game theoretic analysis.⁵⁸ But let us put ourselves in the position of a submarine commander somewhere in the Pacific at the height of the Cold War. A devastating first strike has been launched by the Soviets which has wiped out New York, Chicago, and Los Angeles. With 16 MIRVed missiles in his tubes, the commander receives

⁵⁷ “One of the lamentable principles of human productivity is that it is easier to destroy than create” (Schelling, [43, p. v]).

⁵⁸ At the October 2006 Kennedy School event, Schelling identified one of his principal concerns as “How can you make a promise that's believable when it's clear that left to your own devices you'd rather not do it” [23].

the order to launch a retaliatory strike on Soviet cities, one which, in 20 to 30 min, will kill thirty million people.

Does it make any sense to extract the launch codes, and proceed with firing? What will be accomplished by killing an additional thirty million people? Reputation is irrelevant because this is not a game that will continue. So the submarine commander, if he is rational, cannot logically justify setting in motion the launch procedure, and a rational President cannot logically justify ordering his submarine commander to do so.⁵⁹ But it gets worse.

Since game theory generally presumes that our counterparty is as rational and logical as we are, the U.S. forecasts the likely response of the Soviets to a first strike attack, and concludes that their promise of massive retaliation is hollow—cheap talk. Since if they are rational, they will not retaliate, we can afford to launch first strike against them with impunity.⁶⁰ Moreover, since they will have realized that if the U.S. is rational, our threat of massive retaliation is also hollow, they will have concluded that they can launch with impunity, if we are foolish enough not to have already launched. In each case, the appeal of attack lies both in the prospect of removing a sword of Damocles and in the possibility of territorial aggrandizement or increased world influence. In the absence of any effective defense, there is an enormous game theoretic advantage to the offense, to moving first (Lee, [49, p. 198]).

Schelling understood that no realistic study of international relations, just as no serious study of human behavior, could be premised on the assumption that agents are or should be motivated solely by self-regarding preferences. It is true that such preferences are strong, and underlie what I have called the foraging algorithms (Field, [7]). They cause us to seek food when we are hungry, and water when we are thirsty. But we also possess behavioral predispositions specialized for social interaction, including some that bias us in favor of refraining from first strike, some that make us willing to engage in costly (and thus other-regarding) punishment, even of third party violators, and a heightened sensitivity to detecting those who cheat on the social norms that commonly reflect these species typical behavioral predispositions.

Schelling doesn't address the evolutionary pathways that might have created this. I and others have, and make a strong argument that one must posit the operation of selection at levels higher than

⁵⁹ Some of his advisors felt that President Reagan had doubts he would be able to issue the launch order in such a circumstance, one of the reasons he was so strongly attached to the idea of a missile shield (see Rhodes, [17]).

⁶⁰ Thus Powell does not appreciate that in assuming that “no political objective is worth certain destruction”, the problem of the credibility of US nuclear retaliation to a conventional Soviet attack on Western Europe is the same as the credibility of the threat of nuclear retaliation to a Soviet attack on US territory. And in stating what many would view as self-evident, that with secure second strike, “there is no situation in which it is rational for a state deliberately to launch a nuclear attack first” [53, p. 15]. Powell is simply wrong: this conclusion depends on assuming that one's adversary will in fact retaliate. I am not saying one *should* attack first; if one is rational one will do so only if one believes one's opponent will rationally not retaliate. Irrational human propensities to retaliate are part of the reason in the real world that attacks are deterred. For a similar reason, Jervis is wrong to state “There would be no reason for the Russians to hold back once Americans had destroyed what they value most...” [18, p. 74]. There would be no reason if they were human, but lots of reason if they were rational—since such retaliation would gain them precisely nothing so long as gains are defined in terms of material interest.

the individual organism (group or multilevel selection) in order to allow this to have happened (Wilson and Sober, [50]; Boehm, [51]; Field, [5, 6, 8, 10]; Wilson and Wilson, [52]).

Once one allows for this restraint, once one acknowledges a species typical inhibition on first strike, one which cannot be defended as rational, and one which would not have been favored upon first appearance by organism level selection, one has the foundation for a realistic science of human behavior. One has a framework within which it becomes possible to understand why von Neumann's prediction has, so far, been proved wrong. It becomes possible to see why policies of deterrence and containment can appeal to "natural" human impulses, as much if not more so than the press toward preemption and first strike, which draws support from the logical, rational capabilities associated with the prefrontal cortex. Whereas the counsel of our predispositions specialized for social interaction often agrees with that proffered by our foraging algorithms, sometimes they conflict, and sometimes the former trump or short circuit the latter.

Of course the strength of these inhibitions varies among individuals, for reasons both biological and related to personal history. Of course a skilled leader can defeat them, by demonizing the enemy, or by conjuring threats, both real or imagined. But the point is that the inhibitions are real, they are species typical, and they do have to be defeated (the task of doing so is a central part of military training). It is because of these inhibitions, combined with the real and ever present possibility that they can be defeated by appeal to prudence or rational self-interest, that we have the world we live in, a dangerous world, but one which has been sufficiently peaceful to allow the human population to increase to over 7 billion.

Without these irrational inhibitions, the shadowy world of threats and promises, of the "uglier, more negative, less civilized part of diplomacy" (Schelling, [43], p. vi), would simply not exist. There would be no space for it because conflict would never be looming over us. It always would already have started. But it is the world as it exists, peopled with individuals who have irrational as well as rational behavioral predispositions, that Schelling wished to analyze and understand. That world is one in which the promise or threat of harm becomes a critical element of the grammar of social intercourse, and a central concern of Schelling was how people or nations could be influenced, how these threats or promises are or could be effective. Our world, and the fragile peace in which we live most of the time, presupposes non-Nash behavior. It is a world in which human agents, much but not all of the time, refrain from playing defect even though logic counsels us that defect is the only defensible play.

For most game theorists today, the modeling crucible within which to study the emergence and maintenance of cooperation (non-defection) is the indefinitely repeated game.⁶¹ This is theoretically convenient, because equilibria in which people cooperate (as well as those in which they don't) can, within this context, be attributed to rational self-regarding choice. But it is empirically and historically a poor choice if one is concerned with interactions among agents who have a power to harm each other which is asymmetrically larger than their ability to help each other (aside from failing to harm them). In thinking about behavior among adversaries armed with hydrogen bombs it is absolutely essential that we explain how a Prisoner's Dilemma that might well end up being played

⁶¹ Robert Aumann's contribution to their analysis was a major theme in the announcement of his Nobel Prize. Aumann shared the award with Schelling in 2005.

only once is successfully surmounted, in the sense that neither party defects. To do so one must account for why players are prepared to choose a strategy which in theoretical terms is strictly dominated. Any realistic study of human behavior, whether at the individual, small group, or country level, must begin with the acknowledgment that humans possess some behavioral predispositions that cannot be defended as the rational behavior of a self-regarding agent.

If one or both parties defects in the real time Prisoner's Dilemma, there is no point discussing the logic of deterrence: there is simply no arena for it. Once one enters a state in which a fragile peace is sustained through mutual (and non-rational) restraint on first strike, successful deterrence requires persuading an adversary that you have the intention, under certain states of the world, to behave in a manner that could not at that point be defended as rational. The behavior, and the propensity of humans to indulge in it, is captured well in experiments involving the ultimatum game [54].

The world Schelling (and we) inhabit is a shadowy place where what appears intuitively to make perfect sense cannot in fact be defended as the behavior of a rational agent, where we can talk seriously, if not coherently or consistently, about why it may be advantageous to appear to be irrational. It is a world of conundrums and distorted mirrors, in which we are likely to be startled by our own reflection.

No gas was used by combatants during the Second World War, and to a remarkable degree, the Geneva accords on the treatment of POWs were adhered to by combatants, even though the U.S. never signed these conventions. No nuclear weapons were used in the Korean War, and the U.S. did not bomb across the Yalu river. The Chinese/North Koreans did not bomb bases in Japan, American ships at sea, or bases in Japan (Schelling, [43], pp. 129–130), and Chinese bombers never departed directly from China, always effecting a wheels down in North Korea before pursuing their targets (Schelling, [2]). Thus, the fury on the battlefield notwithstanding, even war has historically been fought *with restraint*, albeit with different amounts of it in different conflicts.

But these restraints, and others like them, always chafe. And within elements of countries such as the United States, there remained and remains a chafing, a chafing at the constraints apparently imposed by a defensive, deterrence based policy, a chafing at the constraints imposed by international diplomacy as symbolized by the United Nations and the various "threats" of world government, a bugaboo that Robert Welch of the John Birch Society railed against. With the breakup of the Soviet Empire and the end of the Cold War, and with the victory of President Bush in the 2000 election, and finally with the unprovoked attack on United States soil on 11 September 2001, voices calling for aggressive preemption were once again front and center, as they had been periodically in the past. In foreign policy, it is true that Iraq was an obsession, but the larger agenda was to redefine our strategic posture.

The Pentagon was to become part of a department of Offense as much as Defense, and the one percent doctrine advanced by Vice President Cheney gave the United States extraordinary scope for launching military action where and when it saw fit. Restraint on first strike was to be drastically attenuated: attacks according to this doctrine could be justified by the merest threat ("one percent probability") of possible attacks on the United States, a threat that would be evaluated and defined by the United States. And along with the attacks on the presumption against the launch of offensive war came attacks on the quaintness of the Geneva conventions, opening the way to the use of

interrogation techniques the U.S. had branded as torture in war crime prosecutions after the Second World War.

As a practical matter the country moved in directions periodically advocated by people like James Burnham, Curtis Lemay, Herman Kahn, John von Neumann, and Barry Goldwater. Science fiction writers and Hollywood screenwriters had fantasized about worlds in which individuals would be incarcerated before they committed crimes, simply because statistical methods predicted a high probability that they would commit offenses. The one percent doctrine represented the application of the principles reflected in the 2002 movie *Minority Report* to countries. As Suskind put it, “Where once a discernible act of aggression against America or its national interest was the threshold for a U.S. military response, now even proof of a threat is too constraining a standard” [47, p. 214]. Although the influence of preempters receded with the evidence of the lack of WMDs in Iraq and the election of Barack Obama, von Neumann’s heirs will surely regain their seat at the policy making table again in the future.

Schelling’s work on deterrence was premised on the assumption that the United States would not itself launch offensive war: that our task rather was to create conditions where we could safeguard our security in a world in which others might. A world of deterrence and containment lacks the simplicity, clarity and dreadful beauty of the Nash equilibrium in the Prisoner’s dilemma. It is a messy world, cluttered with paradoxes, in which arguments can be made on almost all sides of any policy recommendation. But it is the world we must live in if we are to avoid Armageddon. Schelling’s work was not about “thinking the unthinkable” to borrow title words from Herman Kahn’s 1962 book. It was, in part, about avoiding the unthinkable. And to avoid the unthinkable we must have a sound, empirically based picture of the human ethogram, one which acknowledges the sometimes conflicting behavioral predispositions with which we are endowed. People emphasize that we live today in a world of WMDs against which there are only limited defenses. The challenges today are different than they were in the Cold War, but they are not entirely novel. In spite of the hundreds of billions of dollars spent on defense the nation stood completely defenseless against an attack of Soviet ICBMs until the end of the Cold War.

The reality is that in some of the most consequential types of human interaction, formal game theory has not been useful for understanding how people behave, or how they necessarily should behave. Where it is clearly wrong in its predictions, it can, however, serve a useful purpose in helping us break out of the box in which much of modern social science has imprisoned behavioral science. Its usefulness in this fashion becomes apparent when it makes clear, unambiguous predictions which are not borne out by data. In a number of important instances, such as the Prisoner’s Dilemma played once (a game commonly dismissed as “uninteresting” or “too restrictive” by game theorists), its predictions are abundantly contradicted by experimental and observational evidence. But this reality can have less impact than perhaps it should, because some theorists simply aren’t interested in using the theory as a tool for understanding human behavior (Rubinstein, [37]). Deep down, many theorists wish to develop it as a logically consistent and internally coherent set of analyses unconstrained by a requirement that its predictions actually map onto human behavior. Such a posture is not reconcilable with a serious commitment to an empirically based social or behavioral science.

Because of its abstruse, esoteric quality and apparent rigor, many, including economic journalists, continue to be attracted and intrigued by game theory and, more generally rational choice models, without fully appreciating their limitations. Over his long career Schelling did little to discourage those who saw him as a “pioneering strategist” making game theory “serve everyday economics”. But the substance of his work has been premised on acknowledgements of a more complex human ethogram than game theory has been able to accommodate.

If we combine the assertions that deterrence often does work with the argument that it presupposes a commingling of rational and irrational logic and thought processes, then we are forced to question the dominant behavioral theory and motivational assumptions which are thought to underpin realist foreign policy. There are certain conclusions in modern social science which, although indisputably correct, are considered bad form to bring up. A prime example is the absence of an instrumental political rationale for voting in national elections. A second is the inability to defend or explain deterrence when agents are rational and self-regarding. Deterrence works because we are human, not because we are entirely rational. Both of these conclusions point to the limitations of game theoretic approaches. Schelling understood these limitations (even if he did not advertise them) and so should we.

Conflicts of Interest

The author declares no conflicts of interest.

References

1. Warsh, D. *Knowledge and the Wealth of Nations: A Story of Economic Discovery*; Norton: New York, NY, USA, 2006.
2. Schelling, T. Nobel Lecture: An Astonishing Sixty Years: The Legacy of Hiroshima. 2005. Available online: http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2005/schelling-lecture.html (accessed on 11 February 2014).
3. *Handbook of Experimental Economics*; Kagel, J.H., Roth, A.E., Eds.; Princeton University Press: Princeton, NY, USA, 1995.
4. Camerer, C.F. *Behavioral Game Theory: Experiments in Strategic Interaction*; Princeton University Press: Princeton, NJ, USA, 2003.
5. Field, A.J. *Altruistically Inclined? The Behavioral Sciences, Evolutionary Theory, and the Origins of Reciprocity*; University of Michigan Press: Ann Arbor, MI, USA, 2001.
6. Field, A.J. Group Selection and Behavioral Economics. In *Foundations and Extensions of Behavioral Economics: A Handbook*; Altman, M., Ed.; M.E. Sharpe: New York, NY, USA, 2006; pp. 165–182.
7. Field, A.J. Beyond foraging: Behavioral science and the future of institutional economics. *J. Inst. Econ.* **2007**, *3*, 265–291.
8. Field, A.J. Why multilevel selection matters. *J. Bioecon.* **2008**, *10*, 203–238.
9. Thaler, R.H.; Shefrin, H. An economic theory of self-control. *J. Polit. Econ.* **1981**, *89*, 392–406.
10. Field, A.J. Prosociality and the military. *J. Bioecon.* **2014**, *16*, forthcoming.

11. Zagare, F.C.; Kilgour, D.M. *Perfect Deterrence*; Cambridge University Press: Cambridge, UK, 2000.
12. Trachtenberg, M. *History and Strategy*; Princeton University Press: Princeton, NY, USA, 1991.
13. Kahn, H. *Thinking about the Unthinkable*; Horizon Books: New York, NY, USA, 1962.
14. Von Neumann, J.; Morgenstern, O. *Theory of Games and Economic Behavior*; Princeton University Press: Princeton, NY, USA, 1944.
15. Nash, J. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* **1950**, *36*, 48–49.
16. Poundstone, W. *Prisoner's Dilemma*; Doubleday: New York, NY, USA, 1992.
17. Rhodes, R. *Arsenals of Folly: The Making of the Nuclear Arms Race*; Alfred A. Knopf: New York, NY, USA, 2007.
18. Jervis, R. *The Illogic of American Nuclear Strategy*; Cornell University Press: Ithaca, NY, USA, 1984.
19. Royal Swedish Academy of Sciences. Robert Aumann's and Thomas Schelling's Contributions to Game Theory: Analyses of Conflict and Cooperation. 2005. Available online: http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2005/advanced-economicsciences2005.pdf (accessed on 11 February 2014).
20. Schelling, T. *The Strategy of Conflict*; Harvard University Press: Cambridge, MA, USA, 1960.
21. Selten, R. Spieltheoretische behandlung eines oligopolmodells mit nachfragerträge. *Z. Gesamte Staatswiss.* **1965**, *121*, 667–689.
22. Rapoport, A. Review of Schelling (1960). *J. Am. Stat. Assoc.* **1961**, *56*, 433–438.
23. Walt, S.M. Rigor or rigor mortis? Rational choice and security studies. *Int. Secur.* **1999**, *23*, 5–48.
24. Luce, D.; Raiffa, H. *Games and Decisions: Introduction and Critical Survey*; Wiley: New York, NY, USA, 1957.
25. Barkow, J.H.; Cosmides, L.; Tooby, J. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*; Oxford University Press: New York, NY, USA, 1992.
26. Schelling, T. Autobiography. Available online: http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2005/schelling-bio.html (accessed on 11 February 2014).
27. Zagare, F.C. Reconciling rationality with deterrence: A reexamination of the logical foundations of deterrence theory. *J. Theor. Polit.* **2004**, *116*, 107–141.
28. Rosenberg, D.A. The origins of overkill: Nuclear weapons and American strategy, 1945–1960. *Int. Secur.* **1983**, *7*, 3–71.
29. Isaacs, J.; Downing, T. *Cold War: An Illustrated History, 1945–1991*; Little Brown: New York, NY, USA, 1998.
30. Kaplan, F.M. *The Wizards of Armageddon*; Simon and Schuster: New York, NY, USA, 1983.
31. Kaysen, C. Strategic Air Planning and Berlin. Source: National Archives, Record Group 218, Records of the Joint Chiefs of Staff, Records of Maxwell Taylor. 1961. Available online: <http://www.gwu.edu/~nsarchiv/NSAEBB/NSAEBB56/BerlinC1.pdf> (accessed on 11 February 2014).
32. Schelling, T. *Choice and Consequence*; Harvard University Press, Cambridge, MA, USA, 1984.

33. United States National Security Council. 1963. A Study of the Management and Termination of War with the Soviet Union. National Archives, Record Group 59, Department of State Records, Records of Policy Planning Council, 1963–1964, Box 280, File “War Aims”. Available online: <http://www.gwu.edu/~nsarchiv/NSAEBB/NSAEBB31/05-01.htm> (accessed on 11 February 2014).
34. Hobbes, T. *Leviathan, or the Matter, Forme, and Power of a Common-Wealth Ecclesiastical and Civill*; The Clarendon Press: Oxford, UK, 1651.
35. Schelling, T. Review of Rapoport (1964). *Am. Econ. Rev.* **1964**, *54*, 1082–1088.
36. Schelling, T. Comments at A Celebration of the Achievements of Nobel Laureate: Thomas C. Schelling, Kennedy School of Government, Harvard University. 5 October 2006. Available online: <http://forum.iop.harvard.edu/content/celebration-achievements-nobel-laureate-thomas-c-schelling> (accessed on 11 February 2014).
37. Rubinstein, A. A theorist’s view of experiments. *Eur. Econ. Rev.* **2001**, *45*, 615–628.
38. Raiffa, H. *The Art and Science of Negotiation*; Harvard University Press: Cambridge, MA, USA, 1982.
39. Walt, S.M. A model disagreement. *Int. Secur.* **1999**, *24*, 115–130.
40. Rapoport, A. *Strategy and Conscience*; Harper and Row: New York, NY, USA, 1964.
41. Putnam, H. The Place of Facts in a World of Values. In *The Nature of the Physical Universe*; Huff, D., Prewett, O., Eds.; John Wiley: New York, NY, USA, 1979.
42. Field, A.J. Microeconomics, norms, and rationality. *Econ. Dev. Cult. Chang.* **1984**, *32*, 683–711.
43. Schelling, T. *Arms and Influence*; Yale University Press: New Haven, CT, USA, 1966.
44. Lerner, M. *The Age of Overkill: A Preface to World Politics*; Simon and Schuster: New York, NY, USA, 1962.
45. Cheney, D.L.; Seyfarth, R.M. *Baboon Metaphysics: The Evolution of a Social Mind*; University of Chicago Press: Chicago, IL, USA, 2007.
46. Ross, L. The Intuitive Psychologist and His Shortcomings: Distortions in the Attribution Process. In *Advances in Experimental Social Psychology 10*; Berkowitz, L., Ed.; Academic Press: New York, NY, USA, 1977; pp. 173–220.
47. Suskind, R. *The One Percent Doctrine: Deep inside America’s Pursuit of Its Enemies Since 9/11*; Simon and Schuster: New York, NY, USA, 2006.
48. Gintis, H. A framework for the unification of the behavioral sciences. *Behav. Brain Sci.* **2007**, *30*, 1–16.
49. Lee, S.P. *Morality, Prudence, and Nuclear Weapons*; Cambridge University Press: Cambridge, UK, 1993.
50. Wilson, D.S.; Sober, E. Reintroducing group selection to the human behavioral sciences. *Behav. Brain Sci.* **1994**, *17*, 585–654.
51. Boehm, C. *Hierarchy in the Forest: The Evolution of Egalitarian Behavior*; Harvard University Press: Cambridge, MA, USA, 1999.
52. Wilson, D.S.; Wilson, E.O. Rethinking the theoretical foundation of sociobiology. *Q. Rev. Biol.* **2007**, *84*, 327–348.

53. Powell, R. *Nuclear Deterrence Theory: The Search for Credibility*; Cambridge University Press: Cambridge, UK, 1990.
54. Guth, W.; Schmittberger, R.; Schwarze, B. An experimental analysis of ultimatum bargaining. *J. Econ. Behav. Organ.* **1982**, *3*, 367–388.

MDPI AG
Klybeckstrasse 64
4057 Basel, Switzerland
Tel. +41 61 683 77 34
Fax +41 61 302 89 18
<http://www.mdpi.com/>

Games Editorial Office
E-mail: games@mdpi.com
<http://www.mdpi.com/journal/games>



MDPI • Basel • Beijing
ISBN 978-3-03842-014-9
www.mdpi.com

