

Comment: Lawrence E. Blume

The topic of this chapter has been central to the research agendas of Ingela Alger and Jörgen Weibull. Both separately (Lindbeck, Nyberg, and Weibull 1999; Alger 2010) and together (Alger and Weibull 2012), they have made important contributions to the study of prosocial preferences and behavior. The Alger-Weibull research program on the social construction of other-regarding preferences (Alger and Weibull 2013, 2016, this chapter) is exciting both for its formal development of the foundations of evolutionary game theory and for its findings concerning a cultural evolution model of the development of other-regarding preferences. It is conventional to assume that individuals' preferences over social states are concerned only with their own material outcomes. This assumption makes possible the powerful duality between social optimality and market outcomes expressed in the "welfare theorems" and is the baseline environment in which social policy is examined. It is nonetheless naïve in its assumption that preferences are primitives, exogenous in a model of social behavior. The recognition that preferences are to some degree socially constructed challenges many fundamental findings of economic theory; in particular, anything having to do with welfare conclusions. Alger and Weibull have significantly enriched the literature on the social construction of preferences by examining the evolutionary foundations of preference relations.

Evolutionary Game Theory

The Alger and Weibull research program develops new evolutionary game theory tools to say something about the kinds of preferences that would persist in a social system. In the conventional noncooperative theory of

N -person symmetric games, a symmetric Nash equilibrium is a strategy that is a best response for any one individual if it is being used by all the other participants. The fundamental equilibrium concept of evolutionary game theory is that of an evolutionarily stable strategy (ESS).¹ Intuitively, an ESS is a strategy that cannot be “invaded” by another strategy. What does “invasion” mean? Suppose that a large population of individuals are matched at random to participate in the game. An ESS has the property that if a sufficiently large fraction of the population uses it while the remainder of the population uses any other strategy, the expected payoffs to the ESS players are greater than those of the residual population. Evolutionary game theory arose first in biology, and so the strategy alternative to the ESS is said to be invading, and the motivation for the concept is that payoffs measure fitness. Higher payoffs mean higher fitness, and the part of the population using the ESS will have higher average fitness and therefore will outreproduce the group using the invading strategy. To see how the two concepts of Nash equilibrium and ESS fit together, one can check that in any finite and symmetric game, every ESS is a Nash equilibrium of the game. The converse, however, is false.

Alger and Weibull’s program is in the tradition of Güth and Yaari’s (1992) indirect evolutionary approach.² Whereas in traditional evolutionary game theory, preferences are fixed and the evolution of the distribution of actions is governed by the distribution of utility payoffs, in indirect evolutionary models, selection pressure on actions causes the distribution of utilities to evolve. This is analogous to biological models in which selection on phenotypes regulates the distribution of genotypes. In the Güth and Yaari program, behaviors correspond to phenotypes and preferences to genotypes. Payoffs in the game correspond to reproductive fitness.

In the Alger-Weibull program, a strategic interaction is described by a material payoff function π that assigns to each strategy profile a material payoff, (e.g., profit in a model of firm competition).³ Players’ choices are

1. Maynard Smith and Price (1973).

2. See also Güth and Kliemt (1998).

3. Material payoffs, like von Neumann-Morgenstern payoff functions, are linear in the distribution of pure strategy profiles. The present chapter mostly discusses symmetric two-person interactions, but Alger and Weibull (2016) considers multi-player interactions under an aggregative assumption that in the material payoff of an agent’s choices, the choices of others are exchangeable.

governed not by material payoff π , however, but by a payoff function u that represents subjective expected utility preferences over outcomes.

Alger and Weibull repurpose the ESS solution concept from evolutionary game theory as an equilibrium concept for the distribution of von Neumann–Morgenstern payoff functions in the population rather than for the distribution of strategies. In the Alger-Weibull research program, a payoff function u is an ESS if when a sufficiently large fraction of the population uses it while the remainder of the population uses any other payoff function, the expected material payoffs to those with the u payoff function are greater than those of the residual population. A second feature of the Alger and Weibull program—and this is key for their results—is that matching is not random but assortative: Like tends to match with like.⁴

The indirect evolution of preferences with assortative matching produces novel results. The authors label a payoff function a *Homo moralis* payoff function if $u(x, y)$ is an average of $\pi(x, y)$, the material benefit of playing x when others play y , and $\pi(x, x)$, the material benefit when everyone plays according to x . One end of this class is *Homo oeconomicus*, where the averaging weights puts all weight on $\pi(x, y)$, and the other end they label *Homo kantianus*, where all weight is put on the material benefits assuming everyone plays the same.

Other work on the evolution of preferences is close in spirit to the Alger and Weibull program, but different assumptions lead to different outcomes. For instance, Ely and Yilankaya (2001) consider the evolution of preferences in a population using a static stability concept motivated much as is ESS. Because they consider only random matching, they find that outcomes are stable if and only if they are equilibria of the game described by material payoffs; that is, the stable preferences are those of *Homo oeconomicus*. The evolution of social behavior, as opposed to other-regarding preferences, is by now an old topic in evolutionary biology. Hamilton (1964) sees inclusive fitness as an explanation for prosocial behavior, and Grafen (1979) attempts to provide formal support for this idea by considering ESS with nonrandom matching. Bergstrom (1995) considers nonrandom matching for the evolution of altruistic play in a

4. This is not simple to describe in depth, so following the Alger and Weibull essay in this chapter, I shall not attempt to describe it. It is clearly defined in Alger and Weibull (2016, 61).

explicitly biological context and derives *Homo moralis* preferences with $\kappa=1/2$. He called these preferences “semi-Kantian.” There is also some support favoring antisocial preferences. Koçkesen, Ok, and Sethi (2000) introduce a class of payoff functions that depend increasingly on material returns and on relative material returns. Thus if everyone else’s material returns decline while mine do not, then my utility increases. They find that in every equilibrium in a class of games much like those considered by Alger and Weibull but with complete rather than incomplete information, those players with antisocial preferences do materially better than do players who maximize material returns. This is not an evolutionary analysis, but it suggests one. Finally, the Alger and Weibull results work because those with the “right” payoff function receive more material benefits than do others, and sometimes the “right” payoff function is not that of *Homo oeconomicus*. Bester and Güth (1998) and Eshel, Samuelson, and Shaked (1998) develop models where other-regarding preferences do materially worse than does *Homo oeconomicus*, and yet they survive because of group selection effects. The conclusion to draw from this is that details matter for the results of evolutionary models, and we are far from having a complete understanding of how different configurations of environmental characteristics collectively determine evolutionary outcomes. Thus conclusion 3 in the final section of the Alger and Weibull essay in this chapter⁵ is an overstatement. Natural selection does not “favor human motivation in the form of *Homo oeconomicus*.” Different models of natural selection favor different preference relations. *Homo moralis* and *Homo oeconomicus* are two. Nonetheless, Alger and Weibull are to be commended for filling in a new and important part of this landscape.

The promise of the indirect evolutionary approach goes far beyond selection over payoff functions. In evolutionary game theory as received from the biologists, selection forces act on payoffs, and the distribution of strategies evolve. A second level of selection is the indirect evolutionary paradigm. In this case, preferences (which is to say the game itself) evolve in some fashion. Mechanisms for preference evolution include such phenomena as social learning, imitation and other adaptive processes, and

5. p. 410.

the sorting of individuals across roles. These processes operate on a system level rather than at the level of the individual. For instance, Blume and Easley (1992, 2006) show how the redistribution of wealth through repeated trading can drive some kinds of traders from the market. Thus although there are nearly as many behavioral models of choice as there are behavioral economists, only some of them can pass the market survival test. Yet a third level has both strategy choices and the strategic environment coevolve through time. For instance, some papers look at evolution where the community structure, represented by a social network, coevolves with strategic choice and not payoff functions (e.g., Ely 2002; Goyal and Vega-Redondo 2005; Staudigl and Weidenholzer 2014). A novel paper by Sandholm (2002) applies the idea of coevolution of strategies and the game to mechanism design. Moving beyond the Alger and Weibull program, the coevolution of preferences and game forms could contribute much to central questions in political economy in particular and, more generally, the analysis of institutions.

Symmetry

One limitation of their current essay and indeed, the research program, is that Alger and Weibull have so far studied only symmetric environments; that is, for two-player games, those in which the roles of player 1 and player 2 are identical. This limitation is disappointing, because ESS can certainly be generalized to asymmetric games.⁶ Knowing that *Homo moralis* arises in symmetric models—a “one-population” model, one wonders what would emerge from a multiple-population model.

Alger and Weibull consider an asymmetric problem in the third section of their essay in this chapter. I am dissatisfied with their treatment for reasons that foreshadow issues I raise in the sections below having to do with the distinction between positive and normative claims. The strategic situation of their third section imagines a borrower and a lender; the lender has to decide whether to make the loan, and the borrower has to decide whether to pay it back. This is a great example (despite my qualms),

6. See, for instance, Fishman (2008) and citations therein.

because one can see the surprising power of the ESS in preferences. A more conventional analysis would consider repeated interactions between borrowers and lenders. Loans would be made and paid back, because in ongoing relations, reciprocity has value. The borrower understands that if he pays back today, he may be able to get a loan tomorrow. In equilibrium, the lender understands that the borrower understands this, and so she is willing to make the loan. Furthermore, it is her willingness to make future loans that validates the borrower's belief. Alger and Weibull consider only one-shot interactions—there is no possibility of history-dependent behavior. Nonetheless, lending and borrowing can be sustained.

So what is wrong with this? To apply their tools, Alger and Weibull must symmetrize the situation. They state that a canonical way to do this is to initially cover the interaction under a “veil of ignorance” as to who will be in what role. They assume that these roles are contingent. At any moment, a given individual from a single population can either be a borrower or a lender; essentially determined by the flip of a coin. The justification for this move is hinted at by the phrase “veil of ignorance.” They call on the usual suspects—Harsanyi (1953), Rawls (1958), and Vickrey (1945)—who introduced this move in the analysis of social systems. However, the suspects introduced the veil of ignorance, the original position, *ex ante* randomness, for purposes of normative analyses. The original position, behind the veil of ignorance, is a counterfactual hypothetical that provides a frame outside the social system for evaluating the moral consequences of its outcomes. We do not pretend that individuals are actually randomized in such a way. The evolutionary model, however, is concerned with real environments rather than counterfactuals. Of course, there could be situations where roles really are random; a given individual could play one role today and another role tomorrow. But I do not believe that this is a useful way to think about the evolution of preferences where each individual's role is known and certain, set in stone. The use of normative analyses to justify positive claims is one example of the conflation of positive and normative that, I believe, obscures the significance of Alger and Weibull's findings.⁷

7. Their Kantian claims would be much more compelling if a given individual considered the situation of the other party even though she will never ever be in that role. This seems to be required by several of Kant's expressions of his fundamental law.

Welfare Economics

Alger and Weibull have uncovered some powerful results in the positive theory of socially constructed preferences. Their treatment of normative questions, however, and the distinctly normative cast of their entire essay, raises some issues. For instance, how should we view *Homo moralis* preferences from the consequentialist perspective that is traditional in economics? The examples of their third section suggest that a *Homo moralis* world may be materially better than an *Homo oeconomicus* world. To see that this is not the case, consider a variant of the public goods game they discuss. In this variant, N individuals can give, an outcome that is either 0 or 1. The material benefit of the public good is q to each person, the material cost of giving is $0 < c < 1$, and the public good will be provided if and only if the sum of the contributions is at least 1.

Thus, letting $y_{-i} = \sum_{j \neq i} x_j$,

Suppose that $Nq > c$, so that the aggregate material benefit exceeds the cost of provision. It is socially optimal for one individual to provide the good, and the net benefit will be $Nq - c$.

The analysis breaks down into three cases (ignoring boundaries). If $q > c$, one person on his or her own should be willing to give. At one extreme, *Homo kantiensis* chooses to maximize $\pi(x, x)$. The optimum is, $x = 1$, everyone gives, and the public good will be massively oversubscribed. If utilities are interpersonally comparable, the optimum achievable welfare is $Nq - c$, and *Homo kantiensis* society achieves $N(q - c)$, for a material payoff loss of $(N - 1)c$. At the other extreme, *Homo oeconomicus* can achieve efficiency in N distinct asymmetric Nash equilibria. In each equilibrium, one and only one individual gives.

If $q < c$, then *Homo kantiensis* gives zero. The asymmetric pure Nash equilibria of *Homo oeconomicus* also disappear, and *Homo oeconomicus* also gives zero. Both, then, are inefficient.

When $q > c$, there is also a symmetric mixed Nash equilibrium in which the probability of choosing zero is $c^{1/(n-1)}$. In this case, the expected value of the equilibrium to an individual *Homo oeconomicus* is $\rho - c - (\rho - 1)c^{n/(n-1)}$. Comparing this payoff to that of *Homo kantiensis*, we see that it is materially worse when $q > 1$ but materially better when $c/N < q < 1$.

In summary, for $q < c$, both preference types achieve the efficient outcome. For $c < q < 1$, a *Homo kantiensis* society does materially worse than

every Nash equilibrium outcome of a *Homo oeconomicus* society. And for $q > 1$, some *H. oeconomicus* equilibria are efficient, with higher material payoff than that of the *H. kantiensis* society, but the symmetric mixed *oeconomicus* equilibrium is worse.⁸

The general point is that there are problems that, despite being posed symmetrically, have optimal solutions that are asymmetric. Minority games and the related El Farol game provide further examples. This example serves as a caveat to conclusion 4 of the Alger and Weibull essay in this chapter that designing policies for *Homo oeconomicus* when individuals are in fact *Homo moralis* may overincentivize them. Yes, it can, but it may not.

Alger and Weibull's examples in their third section raise the interesting question of how welfare economics should be conducted when preferences are other-regarding. They follow Harsanyi (1980, 1992) and argue that welfare should be measured as the sum of individual material utilities. I followed them in my preceding public good example for purposes of comparison, but this is controversial. To see why, ask: Why exactly is one's desire for a drink of water for herself more necessary to the social welfare calculation than her desire to offer her companion a drink? I can think of two arguments in favor of this claim: one, that water is a necessity for life, and if anything is fundamental, survival needs should be; the other, that to count the companion's welfare in her utility is to double count it. The first argument is nothing more than a statement about marginal rates of substitution at the boundary of the consumption set. At the survival boundary, water for one's self is critical. The second argument says that the utility a decision maker gets from a drink is different from the utility she gets from giving someone else utility. If you take a drink of water, you get some utility. If I offer you that drink, the utility that I get does not count in the social calculation. But if I expend my own resources to do it, the opportunity cost

8. One can derive similar results for the middle-ground cases. The treatment of mixing with *Homo moralis* preferences is unusual, except in the extreme *oeconomicus* case. In Alger and Weibull (2016), we are told that the set X on which $\pi(x, x)$ is defined is the set of mixed strategies in the material game. I understand this to mean that if I were, say, *kantiensis* (just for clarity), and if I chose 1 with probability p , then I assume everyone else is too, and when I consider what happens if I were to choose 1 with probability p' instead, I assume everyone else chooses p' too. This leads to a symmetric randomized equilibrium with an expected social net material benefit that converges upward to $n(p - c)$ as n increases.

of providing the gift does again count. Apparently, only certain actions are allowed to generate utility for welfare purposes. In my view, neither of these arguments holds water.

Alger and Weibull adopt Harsanyi's distinction between personal and social preferences, and they note that one might understand *Homo moralis* as an individual whose personal preferences are the material preferences π and whose social preferences are given by the *Homo moralis* utility function with its degree κ of morality. Ken Arrow famously wrote,⁹ "I am old-fashioned enough to retain David Hume's view that one can never derive 'ought' propositions from 'is' propositions." The findings of evolutionary game theory are "is" propositions. Alger and Weibull are eager to derive from them "oughts." The conflation of "is" and "ought" perhaps undercuts the "is" exercise of their research program.

Alger and Weibull write:

If we take the material payoff function to represent personal utility, then welfare in a society consisting of *Homo moralis* individuals (each with his or her degree of morality) should be defined simply as the sum of their expected material payoffs, just as in ordinary utilitarian welfare theory.

Harsanyi takes personal preferences to be those preferences that guide individuals' choices, their "everyday behavior."¹⁰ If this is what Alger and Weibull mean by personal preferences, then the *moralis* payoff function should represent personal preferences and not material payoffs, and Alger and Weibull's and my welfare calculations are incorrectly done. Harsanyi's description of personal preferences can certainly allow for externalities. If Alger and Weibull believe, following Harsanyi's paradigm, that *moralis* preferences represent what he calls "moral preferences," then I do not understand why they would appear in an evolutionary analysis; decisions are not made based on moral preferences, and so they cannot be selected on.¹¹

If *Homo moralis* preferences are the right preferences to undertake calculations with, then one cannot make welfare comparisons across populations with different degrees κ of morality. By analogy, we might consider two different production economies that differ only in consumers' preferences.

9. In his Ely Lecture, Arrow (1994, 1).

10. Harsanyi (1992, 675).

11. Harsanyi (1992, 671) says that "rational behavior is not a descriptive concept but rather is a normative concept." So he is an odd partner for evolutionary game theory.

We might observe that one economy has a higher GDP than the other, but this gives no guide for comparing the welfare of the two economies, even if utility is interpersonally comparable.

Homo moralis as a Moral Theory

Alger and Weibull write that their work “suggests an evolutionary foundation for a psychologically plausible form of morality, in line with Immanuel Kant’s categorical imperative” (page 392, this chapter). Strictly speaking, they provide “an evolutionary foundation for” preferences that describe behavior consistent with “a psychologically plausible form of morality.” What kind of moral theory? They suggest it is “in line with Immanuel Kant’s categorical imperative.” Bergstrom (1995) uses the phrase “semi-Kantian” to describe *Homo moralis* preferences with $\kappa=1/2$. I believe this Kantian affiliation comes from a misreading of Kant. The idea of Kantian preferences exists outside evolutionary game theory. Roemer (2010) calls a strategy profile in a certain class of games “Kantian” if it is immune to simultaneous proportional deviations from all the players.

Broadly speaking, moral theories fall into one of three classes: consequentialist theories, deontological theories, and virtue theories. Consequentialism emphasizes the consequences of actions. Welfare economics is consequentialist. Deontological theories emphasize duties, rules, and obligations. Most philosophers, including Kant, consider(ed) Kantian theories to be deontological.¹² Virtue ethics emphasizes virtues or moral character. To illustrate, suppose someone’s life is in danger and can be saved by my telling a lie. A consequentialist would lie, because he believes that saving a life is a good outcome. A deontologist would lie if he believed that saving a life when one can without doing injury to others is a universal law. However, if he believed “never lie” is a universal maxim, then he would not lie even to save a life. A virtue ethicist would lie because saving a life is benevolent; a virtue. I claim that *Homo moralis* has much more to do with virtue ethics than with any deontological moral theory.

12. Kagan (2002, 112).

The fundamental moral principle, according to Kant, is a categorical imperative: imperative because it is a command, and categorical because it is required of us unconditionally. That moral principle is, “act only in accordance with that maxim through which you can at the same time will that it become a universal law,” or, in another formulation by Kant, to “act as if the maxim of your action were to become through your will a universal law of nature.” Where does this come from? Kant wrote:¹³

Everyone must admit that a law, if it is to be valid morally, i.e., as the ground of an obligation, has to carry absolute necessity with it; that the command “You ought not to lie” is valid not merely for human beings, as though other rational beings did not have to heed it; and likewise all the other genuinely moral laws; hence that the ground of obligation here is to be sought not in the nature of the human being or the circumstances of the world in which he is placed, but a priori solely in concepts of pure reason, and that every other precept grounded on principles of mere experience, and even a precept that is universal in a certain aspect, insofar as it is supported in the smallest part on empirical grounds, perhaps only as to its motive, can be called a practical rule, but never a moral law.

In other words, it is to be rationally derivable, assuming that every human were to heed it. The law is based entirely on reason and is not a consequence of any facts on the ground. In particular, moral propositions are to be independent of whom they are applied to; their preferences make no difference. These propositions are independent of our desires and uncoupled from the consequences that ensue. Clearly, however, the rules that one would derive from *Homo moralis* preferences depend on what the material payoffs are: Consequences matter. To put this somewhat differently, Kant’s categorical imperative has a game-theoretic nature: An assumption about the behavior of others enters into your calculation about how you should behave. But *Homo kantiansis* is not Kantian, because his evaluation of the act is independent of his preferences. If a given maxim survives the categorical imperative test, one is obliged to act according to it, even if it is preference minimal. Thus the moral theory for which Alger and Weibull “provide an evolutionary foundation” is not Kantian. Quite the opposite. Harsanyi (1980) calls individuals who maximize a class of utility functions

13. Kant ([1785] 2002, 5).

containing *Homo moralis* payoff functions rule “utilitarians.” It appears to be consequentialist.¹⁴

To the extent that we use the language of choice theory to talk about moral choices of individuals, any such theory will appear to be consequentialist. One can read virtue ethics this way. Our preferences are shaped by our character. Thus in some situations, preferences of individuals who have internalized particular virtues will look different than those of individuals who have not. And so the choices of those of virtuous character—sympathetic, charitable, etc.—will reflect these virtues. These are moral choices. One school of modern virtue ethics, so-called agent-based ethics,¹⁵ “understands rightness in terms of good motivations and wrongness in terms of the having of bad (or insufficiently good) motives.” Alger and Weibull’s evolutionary account of preference evolution supports this view. They tell us that, as a consequence of the social condition, as a result of social interaction, preferences must in the long run take on a certain form, and that form is other-regarding.

Adam Smith ([1759] 2004, 1) begins *The Theory of Moral Sentiments* by claiming the universality of certain virtues:

How selfish soever man may be supposed, there are evidently some principles in his nature, which interest him in the fortune of others, and render their happiness necessary to him, though he derives nothing from it except the pleasure of seeing it. Of this kind is pity or compassion, the emotion which we feel for the misery of others, when we either see it, or are made to conceive it in a very lively manner. That we often derive sorrow from the sorrow of others, is a matter of fact too obvious to require any instances to prove it; for this sentiment, like all the other original passions of human nature, is by no means confined to the virtuous and humane, though they perhaps may feel it with the most exquisite sensibility. The greatest ruffian, the most hardened violator of the laws of society, is not altogether without it.

He goes on to argue¹⁶

that this is the source of our fellow-feeling for the misery of others, that it is by changing places in fancy with the sufferer, that we come either to conceive or to

14. In fairness, I should say that the contrast between deontology and consequentialism is not as sharp as it is often made out to be and is somewhat contested. See Kagan (2002) and Cummiskey (1990).

15. Slote (2001, 14).

16. Smith ([1759] 2004, 4).

be affected by what he feels, may be demonstrated by many obvious observations, if it should not be thought sufficiently evident of itself.

This expression of sympathy is, for Smith, the source of our moral decision-making. In a passage that is reminiscent of *Homo moralis*, he states:¹⁷

The principle by which we naturally either approve or disapprove of our own conduct, seems to be altogether the same with that by which we exercise the like judgments concerning the conduct of other people. We either approve or disapprove of the conduct of another man according as we feel that, when we bring his case home to ourselves, we either can or cannot entirely sympathize with the sentiments and motives which directed it.

Assuming others behave as x , how do we feel about x ?

Finally, it is interesting to note that perhaps Smith in the *Theory of Moral Sentiments* would be sympathetic to the Alger and Weibull program. He writes:¹⁸

It is thus that the general rules of morality are formed. They are ultimately founded upon experience of what, in particular instances, our moral faculties, our natural sense of merit and propriety, approve, or disapprove of. We do not originally approve or condemn particular actions; because, upon examination, they appear to be agreeable or inconsistent with a certain general rule. The general rule, on the contrary, is formed, by finding from experience, that all actions of a certain kind, or circumstanced in a certain manner, are approved or disapproved of.

Our moral views emerge from experience, a social process. It would be asking too much of the mid-seventeenth-century Smith to distinguish between social learning and social evolution, and even today, it is not clear that, as classes, these are observationally distinct. But Alger and Weibull need not commit to a mechanism for their ESS analysis beyond the fact that it is monotone in payoffs, and so they are not inconsistent with Smith.

Conclusion

Although I have reservations about Alger and Weibull's (and many other economists) assertions about moral theory, the Alger and Weibull research program is among the most ambitious and promising to date on the exploration of the evolution of other-regarding preferences. The results are

17. Smith ([1759] 2004, 151–152).

18. Smith ([1759] 2004, 206).

exciting both for what they find and for the extent of the environments in which they hold.¹⁹ Received game and market theory is of the take-all-comer's variety; equilibrium exists no matter what preferences agents hold. But if preferences are socially constructed, the forces described by Alger and Weibull should limit the kinds of preferences that are prevalent. Game and market theory should take advantage of this fact to make sharper predictions about the behavior of social systems. Finally, Jörgen Weibull has contributed significantly to the literature on evolutionary dynamics, and so I look forward to seeing this program progress from the static analysis of ESS to the much harder (but potentially richer) dynamic analyses that have emerged in recent years.

References

- Alger, Ingela. 2010. "Public Goods Games, Altruism, and Evolution." *Journal of Public Economic Theory* 12 (4): 789–813.
- Alger, Ingela, and Jörgen W. Weibull. 2012. "A Generalization of Hamilton's Rule—Love Others How Much?" *Journal of Theoretical Biology, Evolution of Cooperation* 299: 42–54.
- Alger, Ingela, and Jörgen W. Weibull. 2013. "Homo moralis—Preference Evolution under Incomplete Information and Assortative Matching." *Econometrica* 81 (6): 2269–2302.
- Alger, Ingela, and Jörgen W. Weibull. 2016. "Evolution and Kantian Morality." *Games and Economic Behavior* 98: 56–67.
- Arrow, Kenneth. 1994. "Methodological Individualism and Social Knowledge." *American Economic Review* 84 (2): 1–9.
- Bergstrom, Theodore C. 1995. "On the Evolution of Altruistic Ethical Rules for Siblings." *American Economic Review* 85 (1): 58–81.
- Bester, Helmut, and Werner Güth. 1998. "Is Altruism Evolutionarily Stable?" *Journal of Economic Behavior and Organization* 34 (2): 193–209.
- Blume, Lawrence, and David Easley. 1992. "Evolution and Market Behavior." *Journal of Economic Theory* 58 (1): 9–40.

19. This is not apparent in the present chapter, but can be seen in Alger and Weibull (2016).

- Blume, Lawrence, and David Easley. 2006. "If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets." *Econometrica* 74 (4): 929–966.
- Cummiskey, David. 1990. "Kantian Consequentialism." *Ethics* 100 (3): 586–615.
- Ely, Jeffrey C. 2002. "Local Conventions." *Advances in Theoretical Economics* 2 (1): 1–32.
- Ely, Jeffrey C., and Okan Yilankaya. 2001. "Nash Equilibrium and the Evolution of Preferences." *Journal of Economic Theory* 97 (2): 255–272.
- Eshel, Ilan, Larry Samuelson, and Avner Shaked. 1998. "Altruists, Egoists, and Hooligans in a Local Interaction Model." *American Economic Review* 88 (1): 157–179.
- Fishman, Michael A. 2008. "Asymmetric Evolutionary Games with Non-linear Pure Strategy Payoffs." *Games and Economic Behavior* 63 (1): 77–90.
- Goyal, Sanjeev, and Fernando Vega-Redondo. 2005. "Network Formation and Social Coordination." *Games and Economic Behavior* 50 (2): 178–207.
- Grafen, Alan. 1979. "The Hawk-Dove Game Played between Relatives." *Animal Behaviour* 27 (3): 905–907.
- Güth, Werner, and Hartmut Kliemt. 1998. "The Indirect Evolutionary Approach: Bridging the Gap between Rationality and Adaptation." *Rationality and Society* 10 (3): 377–399.
- Güth, Werner, and Menahem E. Yaari. 1992. "Explaining Reciprocal Behavior in Simple Strategic Games: An Evolutionary Approach." In *Explaining Process and Change: Approaches to Evolutionary Economics*, edited by Ulrich Witt, 23–34. Ann Arbor: University of Michigan Press.
- Hamilton, William D. 1964. "The Genetical Evolution of Social Behavior, I and II." *Journal of Theoretical Biology* 7 (1): 1–52.
- Harsanyi, John C. 1953. "Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking." *Journal of Political Economy* 61 (5): 434–435.
- Harsanyi, John C. 1980. "Rule Utilitarianism, Rights, Obligations and the Theory of Rational Behavior." *Theory and Decision* 12 (2): 115–133.
- Harsanyi, John C. 1992. "Game and Decision Theoretic Models in Ethics." In *Handbook of Game Theory with Economic Applications*, volume 1, edited by Robert J. Aumann and Sergiu Hart, 669–707. Amsterdam: North-Holland.
- Kagan, Shelly. 2002. "Kantianism for Consequentialists." In *Groundwork for the Metaphysics of Morals*, edited by Allen W. Wood, 111–156. New Haven, CT: Yale University Press.
- Kant, Immanuel. [1785] 2002. *Groundwork for the Metaphysics of Morals*. Edited and translated by Allen W. Wood. New Haven, CT: Yale University Press.

Koçkesen, Levent, Efe A. Ok, and Rajiv Sethi. 2000. "Evolution of Interdependent Preferences in Aggregative Games." *Games and Economic Behavior* 31 (2): 303–310.

Lindbeck, Assar, Sten Nyberg, and Jörgen W. Weibull. 1999. "Social Norms and Economic Incentives in the Welfare State." *Quarterly Journal of Economics* 114 (1): 1–35.

Maynard Smith, John, and George R. Price. 1973. "The Logic of Animal Conflict." *Nature* 246 (November): 15–18.

Rawls, John. 1958. "Justice as Fairness." *Philosophical Review* 67 (2): 164–194.

Roemer, John E. 2010. "Kantian Equilibrium." *Scandinavian Journal of Economics* 112 (1): 1–24.

Sandholm, William H. 2002. "Evolutionary Implementation and Congestion Pricing." *Review of Economic Studies* 69 (3): 667–689.

Slote, Michael. 2001. *Morals from Motives*. Oxford: Oxford University Press.

Smith, Adam. [1759] 2004. *The Theory of Moral Sentiments*. New York: Barnes and Noble.

Staudigl, Mathias, and Simon Weidenholzer. 2014. "Constrained Interactions and Social Coordination." *Journal of Economic Theory* 152: 41–63.

Vickrey, William. 1945. "Measuring Marginal Utility by Reactions to Risk." *Econometrica* 13 (4): 319–333.

This is a section of [doi:10.7551/mitpress/11130.001.0001](https://doi.org/10.7551/mitpress/11130.001.0001)

The State of Economics, the State of the World

Edited by: Kaushik Basu, David Rosenblatt,
Claudia Sepúlveda

Citation:

The State of Economics, the State of the World

Edited by: Kaushik Basu, David Rosenblatt, Claudia Sepúlveda

DOI: 10.7551/mitpress/11130.001.0001

ISBN (electronic): 9780262353472

Publisher: The MIT Press

Published: 2020



The MIT Press



This work is available under the Creative Commons Attribution—NonCommercial—NoDerivatives 3.0 IGO license (CC BY-NC-ND 3.0 IGO) <http://creativecommons.org/licenses/by-nc-nd/3.0/igo>.

Some rights reserved

The findings, interpretations, and conclusions expressed in this work are those of the authors and do not necessarily reflect the views of The World Bank, its Board of Executive Directors, or the governments they represent. The World Bank does not guarantee the accuracy, completeness, or currency of the data included in this work and does not assume responsibility for any errors, omissions, or discrepancies in the information, or liability with respect to the use of or failure to use the information, methods, processes, or conclusions set forth. The boundaries, colors, denominations, and other information shown on any map in this work do not imply any judgment on the part of The World Bank concerning the legal status of any territory or the endorsement or acceptance of such boundaries.

Nothing herein shall constitute or be construed or considered to be a limitation upon or waiver of the privileges and immunities of The World Bank, all of which are specifically reserved.

Attribution—Please cite the work as follows: The World Bank. 2019. *The state of economics, the state of the world* / edited by Kaushik Basu, Claudia Sepulveda, and David Rosenblatt. Published by MIT Press. © World Bank. License: Creative Commons Attribution—NonCommercial—NoDerivatives 3.0 IGO (CC BY-NC-ND 3.0 IGO).

Third-party content—The World Bank does not necessarily own each component of the content contained within the work. The World Bank therefore does not warrant that the use of any third-party-owned individual component or part contained in the work will not infringe on the rights of those third parties. The risk of claims resulting from such infringement rests solely with you. If you wish to re-use a component of the work, it is your responsibility to determine whether permission is needed for that re-use and to obtain permission from the copyright owner. Examples of components can include, but are not limited to, tables, figures, or images.

All queries on rights and licenses should be addressed to the Publishing and Knowledge Division, The World Bank, 1818 H Street NW, Washington, DC 20433, USA; fax: 202-522-2625; e-mail: pubrights@worldbank.org.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Basu, Kaushik, editor. | Sepúlveda, Claudia Paz, 1969– editor. | Rosenblatt, David, editor.

Title: *The state of economics, the state of the world* / edited by Kaushik Basu, Claudia Sepulveda, and David Rosenblatt.

Description: Cambridge, MA : MIT Press, [2019] | Includes bibliographical references and index.

Identifiers: LCCN 2018046336 | ISBN 9780262039994 (hardcover : alk. paper)

Subjects: LCSH: Economic development. | Information technology—Economic aspects. | Monetary policy. | Social change.

Classification: LCC HD82 .S8223 2019 | DDC 330.1—dc23

LC record available at <https://lcn.loc.gov/2018046336>

10 9 8 7 6 5 4 3 2 1