Stable Belief Structures: Social Stratification and Nash Equilibrium

By Roger A. McCain

> "Of course, the decision-maker may also hold inaccurate beliefs. She may, for example, overestimate the prevalence of some group traits…."
>
> Rose, 2023, p. 906

> "Our findings shed light on the role of beliefs about race, as opposed to racial animus, in contributing to racial differentials in well-being."
>
> Shari et al, 2023, p. 924

Inequality of wealth, status, and income seem to be quite persistent (Björklund and Jäntti 1997, Clark et al. 2012, Clark and Cummins 2014, 2015, Guell et al. 2015, Torche, 2015, Long and Ferrie 2013). While the persistence of socially ascribed status, such as caste or race, does not seem to have been as extensively studied, informal historical information suggests that it is no less and perhaps more persistent. Consider, for example, the status differences created by the Norman conquest of England. Clark et al. find that, almost 1000 years later, people who share the conquerors' surnames are statistically overrepresented in the enrollment at Oxford university, with a frequency three times as great as their frequency in the general population. And this persistence itself persists into the periods of industrialization and the rise of social democracy.

John Nash offered his noncooperative equilibrium concept (Gibbons, 1997, McCain 2023, Ch. 4) in game theory as a model of interactive decisions that would be stable as customs are stable. Indeed, some consistent patterns of decisions, such as driving on the left or right side of the road, combine a seemingly arbitrary status with remarkable stability, and comprise equilibria of appropriately configured games (McCain, 2023, pp. 91-93). Caste systems and similar systems of social stratification similarly combine arbitrary and at best unproductive status differences with great persistence. Perhaps the roots of this persistence may be found in some structure such as Nash equilibrium. It is true that noncooperative game theory makes a number of simplifying assumptions about such things as the agents' common knowledge of the structure of the "game" and the payoffs or preferences of other agents. However, there is now a well-established body of theory that relaxes some of those simplifying assumptions, Bayesian Nash equilibrium theory[1], in which some agents may be unaware ex ante of some of the payoffs and strategies of other agents, and form opinions on those matters only on the basis of experience on the "play of the game." In such a case, beliefs that are objectively false may be consistent with experience and thus very persistent[2]. This paper will present some examples to explore the implications of this point for social stratification[3].

In a Bayesian Nash equilibrium, agents may be of different "types" with different preferences or payoffs and possibly different capabilities. When an agent is matched to play a "game," the agent may be unaware of the types of other agents in the game, and able to decide on the agent's own course of action only on the basis of probability judgments. In many applications of Bayesian Nash equilibrium, there are only two information states, First, each agent knows their own type with probability one. Second, each agent judges the probability of the type of another agent according to common prior probabilities, based on all commonly available information. In such an example there are no signals of type. By contrast, in other applications, decision makers may be alert for signals that are correlated with the type of the counterparts in a particular match-up: for example, reputation may be an important signal.

i. Effort Games

Consider Game 1, shown in Table 1. Two agents, 1 and 2, collaborate in some productive task that yields some reward to each of them. They each choose between two strategies: participate without any particular effort, coded as "shirk," or complete the joint task with a great but unpleasant effort, coded as "work." The numbers in the cells are the gross payoffs to the

---

[1] Harsanyi, 1967. Gibbons, 1997, introduces Bayesian equilibrium in a general survey of game theory. McCain 2023, pp. 146-151 gives a further simple example.

[2] Some recent studies have addressed the formation of opinions with models in the tradition of economic and game theory. Compare the ideas described here with Nunn, 2022, Lyonnet and Stern 2023.)

[3] Compare Garcia and Darity, Jr. (2022) for a broadly similar game-theoretic discussion of decisions in the presence of a racial dichotomy and draws some similar conclusions. Their model uses more conventional, rather than Bayesian Nash equilibrium.

agents on an arbitrary scale of 10. However, great effort is unpleasant so that, when an agent makes a great effort, the discomfort is equivalent to the deduction of an effort penalty ε.

Table 1. Game 1, An Effort Dilemma

		2	
		work	shirk
1	work	$10-\varepsilon,10-\varepsilon$	$6-\varepsilon,9$
	shirk	$9,6-\varepsilon$	6,6

Clearly, a great deal will depend on the value of ε. Suppose $2>\varepsilon>1$. Then "shirk" is the dominant strategy, but "work, work" yields at least 8,8, while the play of the dominant strategy yields 6,6. This, of course, corresponds with the famous Prisoner' Dilemma. It is often taken as a description of team production in general and as the basis of an argument that effort commitment needs some enforcement in team production. Like other examples similar to the Prisoners' Dilemma it illustrates how perverse incentives may make it difficult to arrive at a mutually beneficial outcome.

Suppose, however, that $0<\varepsilon<1$. Then there are two Nash equilibria, one when both choose "work" and one when both choose "shirk." The problem then is not perverse incentives but a lack of information, in that, to choose a best response, each agent must conjecture as to which strategy the other will choose. The "work, work" equilibrium is payoff dominant, however, and while "shirk, shirk" is risk dominant, there are no loss outcomes in this game. Thus it is relatively plausible that rational agents would choose the payoff-dominant equilibrium, as Harsanyi and Selten argue (1988). However, there might be some common signal from outside the game, a "sunspot," to determine a Schelling focal point, and that might stabilize a "shirk, shirk" equilibrium.

These examples assume that agents are averse to effort, and that all are effort averse to an extent that their aversion can be measured by the same penalty. But it is likely that individuals differ with respect to their degree of effort aversion, and there is no reason to suppose that (in the context of Game 1) 1 is a lower limit for the effort penalty. In a large population we might expect to find individual values of ε over a wide range of nonnegative values. However, for simplicity, we will assume that all agents are of one of two types. For type A, ε=0.5, while for type B, ε=1.5. Then envision a large population of both types, who are then matched at random in each period to play Game 1. In a particular match, suppose Agent 1 is of type A, and Agent 1 believes[4] that Agent 2, with whom Agent 1 is matched at this round, is of type B. Then Agent 1 believes that the payoffs will be as in Game 2.

Table 2. Game 2

		2	
		work	shirk
1	work	9.5,8.5	5.5,9
	shirk	9,4.5	6,6

If Agent 1's assumptions are correct, then Agent 2 has a dominant strategy: shirk. Then, by elimination of dominated strategies, "shirk, shirk" is the only Nash equilibrium. Suppose, further, that Agent 2 is in fact of type A, and that Agent 2 knows that Agent 1 believes that Agent 2 is of type B. Then Agent 2 is aware that Agent 1 will choose "shirk," so Agent 2's best response is "shirk." Agent 1's mistaken belief that Agent 2 is of type B has made the mutually

---

[4] Of course we cannot say that Agent 1 knows that Agent 2 is of type B, since one cannot know anything that is not true: knowledge requires "justified true belief" at the least. (Ichikawa and Steup, 2018). However, clearly it is possible to believe what is not true. Here we can do no better than to follow Aumann, 199, p. 304, expression 12.2, and note that while one will assign a probability of one to a proposition the person knows or to an event whose occurrence is certain, one may assign a probability of one to other propositions or events. That is the meaning ascribed to "belief" in this paper.

beneficial “work,” “work” outcome unattainable. But, reflecting on his experience, Agent 1 will see that he has been proven right: Agent 2 has chosen “shirk,” just as Agent 1 predicted.

What this illustrates is that beliefs may be persistent, even though contrary to fact, because they are founded on experience that is stable as it corresponds to Nash equilibrium in common interactive occasions[5]. In this example, there might be no type B agents at all, but the belief that there are some and that they can be recognized generates decisions that seem to confirm the belief.

If in fact there are some agents of two or more distinct types, agents may adjust their decisions according to a mathematical expectation of payoffs, with probability judgments as to the types others may be. In many applications there are just two probability states: 1) each person knows their own type with probability 1 and 2) the probabilities of types of others correspond to commonly held prior probabilities (Aumann, R. J. 1976). However, a reasonable player might be alert for signals correlated with the types of others. Reputation could be one such.

Again, then, suppose that agents are matched at random in each period to play Game 1, but that some reliable record exists of each agent’s past strategy decisions. If Agent 1 is of type A and observes that in the past Agent 2 has chosen “work” consistently, Agent 1 may infer that Agent 2 is of type A also, so that “work, work” is a Nash equilibrium. But if Agent 2 is in fact of type B, Agent 2’s rational choice may be to choose “work,” in order to maintain the agent’s reputation as being of type A, so that the agent can look forward to a series of payoffs of 8.5, rather than 6, in future matchups. The mathematics of this decision is similar to the model of repeated play of a social dilemma between the same two agents for an uncertain number of periods, although the mechanism is different. We can envision an equilibrium in which it is believed that all or most agents are of type A, when in fact none are. Instead, all or most choose the strategy “work” to create or maintain their reputations as type A. But the result is that that it is widely and falsely believed that most agents are of type A.

This example further illustrates how mistaken beliefs may be persistent because they are consistent with Nash equilibria, but it has two special characteristics. First, information about past decisions is available and reliable. Second, the agent has a choice whether to build a reputation or not: the label as a worker is chosen, not attributed. But, while a reputational

---

[5] Compare Dorsinville, Jean L. (2023); contrast Bordalo, et al. (2016), Nunn (2022).

equilibrium could be stable, it is not clear how the system might make a transition to it. If it were widely believed that most agents are of type B, both kinds of agents, applying a common prior probability that the counterpart agent is of type B, would choose "shirk," and there would be nothing to be gained by building a reputation to the contrary.

But, for the purposes of this paper, we are more concerned with signals that the agent does not choose but are attributed regardless of the agent's decision. Suppose, then, that information to verify reputations is not available, but that agents are "labelled" in various ways. In McCain 2018 p. 251, I gave the example of big ears. In the actual world skin color, dialects or languages, and surnames (Clark and Cummins, 2014, 2015, Guell, Rodriguez Mora, and Tellmer 2015, Shaffer 2004, Shari et al. 2023) might be easily observed labels. Suppose, then, that the population has two components, some of whom bear some visible label while others do not. Suppose further that, among the unlabeled[6] population, it is widely believed that labelled agents are of type B, and most labelled agents know that many unlabeled agents believe this. Then matches of labeled with unlabeled agents will usually fail to realize the "work, work" equilibrium, and the Nash equilibrium realized will be consistent with the widely held, but mistaken, belief on the part of unlabeled agents that labeled agents are of type B. The example of big ears in the 2018 book was precisely to make the point that the equilibrium is only one of several, and might not be realized in the actual world: that there is nothing natural or inevitable about it.

Now consider Game 3, a three-person effort game. Once again, if $\varepsilon<1$, there may be two equilibria, with the work equilibrium Pareto-preferable to the other. Suppose that in fact all three agents are of type A with $\varepsilon<1$. Suppose, however, that Agent 2 is labeled and that consequently Agent 1 is certain that Agent 2 is of type B with $\varepsilon>1$. Suppose, further, what while Agent 3 does not believe that Agent 2 is of type B, Agent 3 knows with certainty that Agent 1 believes this, and that Agent 2 knows that Agent 1 believes this. Then Agent 3 reasons that Agents 1 and 2 will choose "shirk," so then Agent 3's best response is to choose "shirk," even though Agent 3 does

---

[6] While the meaning of this term will be clear, the term itself is logically contradictory, since the lack of a "label" will in effect serve as a label in the example. It will be used here in the interest of brevity.

not suppose that Agent 2 is of type B. Thus, again, the beliefs of some agents lead to a Nash equilibrium that is consistent with those beliefs, even though the beliefs are mistaken.

Table 3. Game 3

		3			
		Work		Shirk	
		2		2	
		work	shirk	work	shirk
1	work	10-ε, 10-ε,10-ε	6-ε,9,6-ε	6-ε,6-ε,9	6-ε,6,6
	shirk	9, 6-ε,6-ε	6, 6, 6-ε	6, 6-ε,6	6,6,6

To make Game 3 a little more concrete, we might suppose that Agents 1 and 2 are entrepreneurs and Agent 3 is a banker. Agent 2 is seeking a loan with a business plan that relies on a B to B service that may or may not be supplied by Agent 1. The banker, believing that because of his prejudice Agent 1 will not collaborate with Agent 2, refuses to make the loan to Agent 2. (See e.g. Blanchflower et al. 2003, Garcia and Darity, 2022. The reflections in Dorsinville, Jean L, 2023 are suggestive.)

ii. Matching in Game 3

This section will explore the implications of mistaken beliefs about type for a matching game to play the three-person effort game, Game 3. As before, for type A, ε=0.5 and for type B, ε=1.5. Suppose the proportion of agents of type A is the same in both labelled and unlabeled populations and is q>0.354=SQRT(0.125). Agents who hold the objectively false belief that all labelled persons are of type B will be referred to as believers. The proportion of the population who are labelled is p and the proportion of unlabeled who are believers is r.

Case 1. The probability that a random match would group three unlabeled agents is $(1\text{-}p)^3$. The probability that all are type A is $(1\text{-}q)^3$. Take the viewpoint of a type A agent in such a match. Can "work" be a best response for this agent? We seek a rule for strategic responses for type A agents in such a matchup: either all type A agents choose "work" or "shirk." The agent knows his own type and that the probability that the other two are both of type A is q^2. If they are, then the payoff to "work" is 9.5; otherwise 5.5. Then the expected value of choosing according to the rule is

1. $EV(1) = q^2(9.5) + (1 - q^2)(5.5) = 5.5 + 4q^2$

while a rule of "shirk" has an expected value of 6. Since by assumption $q>0.354$, *EV(1)*>6, and thus the rule leads to better results than the alternative rule, and "work" may be a best response in this case.

Case 2. The probability that a random match would group three labeled agents is p^3. The same reasoning then applies.

Case 3. The probability that two unlabeled agents would be matched with a labeled agent is $(1\text{-}p)^2p$. If both unlabeled agents are believers, with probability r^2, or if just one is a believer with probability (1-r)r, then the "all work" equilibrium cannot occur. Suppose member 1 is an unlabeled unbeliever of type A. The probability that others are of type A is q^2 and the probability that the second unlabeled agent is an unbeliever is (1-r). Then the expected value of the "work" strategy is

2. $EV(3) = q^2(1 - r)(9.5) + \left(1 - q^2(1 - r)\right)(5.5) = 5.5 + 4(1 - r)q^2$

In this case, "work" is a best response only if $q^2(1 - r) > 0.125$. This threshold requires that q be greater than would be greater than would suffice in Cases 1 and 2. In other words, depending on the frequency of type A, it may be impossible for a mixed match of this kind to realize an "all work" equilibrium even though all are of type A. The probability of this outcome also increases with the proportion of unlabeled persons who are believers.

Case 4. The probability that two labeled agents would be matched with one unlabeled agent is $p^2(1\text{-}p)$. Take the point of view of a labeled agent of type A. The probability that the other two agents are both of type A is q^2 and the probability that the unlabeled agent is an unbeliever is (1-r). Again, the expected value of "work" is as shown in equation 2, and the same result follows.

We may draw three inferences from this discussion. First, even in matches without any believers, all of whom are type A, the presence of believers in the population reduces the likelihood that mixed groups will arrive at the productive "all work" equilibrium. Second, this does not apply to matches that are unmixed, either all labeled or all unlabeled. Finally, to the extent that agents may be able to reject a matching and expect to find another, this may give type A agents, labeled or unlabeled, reason to reject a mixed matching.

iii. Formalization

This section will sketch a formalization that generalizes the examples in sections i and ii. The examples given are special cases of this more general framework, but special cases are sufficient for proofs of existence, which is the objective of this paper.

A game Γ comprises

1) A set $\mathbb{N} = \{1,2,\dots,n\}$ of decision-makers,
2) A set $\mathbb{T} = \{i, ii, \dots, m\}$ of types, where m≤n,
3) A set $\Omega = \{\alpha, \beta, \dots, \zeta\}$ of outcomes,
4) A map h=H(i) with $h \in \mathbb{T},\ i \in \mathbb{N}$, such that h=H(i) ⟺i is of type h,
5) For each type y a set $S_y = \left\{\sigma_{1y}, \sigma_{2y} \dots, \sigma_{q_y,y}\right\}$ of strategies and a preference system $\mathcal{P}_y$.

Remark: In much of the literature of game theory, and in the examples shown above, items 3) and 5) are simplified as follows: outcomes are supposed to be numerical vectors of dimension n, with the ith element understood as the "payoff" to agent i. Then the preferences are assumed to be such that larger "payoffs" are preferred to smaller. It is often further assumed that preferences are regular enough that preferences can be defined over probability distributions of outcomes, or further that decision-makers are risk-neutral. Risk-neutrality is assumed in some of the discussion in the previous sections.

6) An outcome function $\omega = F(\sigma_1, \sigma_2, \dots, \sigma_j, \dots, \sigma_n)$ where h=H(i) ⟹ $\sigma_i \in S_h$, $\omega \in \Omega$.

Remark: This says that if agent i is of type h, then strategy σ_i must be a member of the set of strategies available to agents of type h. In much game theory (other than Bayesian Nash equilibrium theory) this is simplified either by assuming that all agents are of the same type or that each is of a different type, but types may be known to all. For the examples in this paper, both types have the same strategies available, but different types differ only in their preferences. This is a fairly common assumption in Bayesian Nash equilibrium theory.

7) For each agent i a belief system expressed as a probability distribution $p_{i,y,j} = \psi_i(y,j)$, where $i \in \mathbb{N}, j \in \mathbb{N}$, j≠i, $y_{i,j} \in \mathbb{T}$. Here $p_{i,y,j}$ is agent i's estimate of the probability that agent j is of type y.
8) For each agent i a revision function $\phi_i(y,j) = \Psi_i(\omega, \psi_i(y,j))$, where $\phi_i(y,j)$ is agent i's belief system revised in the light of the actual occurrence of outcome ω. $\phi_i(y,j)$ might be arrived at by an application of Bayes' rule, if we assume that agents are rational in a certain sense. If ω is the expected outcome, then presumably $\Psi_i(\omega, \psi_i(y,j))$=$\psi_i(y,j)$.

Suppose that, for each $i \in \mathbb{N}$, i of type y, we have strategies $\sigma_{i,l} \in S_y$ such that

a) $\sigma_{i,k} \neq \sigma_{i,l}, \sigma_{i,k} \in S_y \Rightarrow \sim F(\sigma_1, \sigma_2, \dots, \sigma_{i,k}, \dots, \sigma_n)\mathcal{P}_i F(\sigma_1, \sigma_2, \dots, \sigma_{i,l}, \dots, \sigma_n)$

That is, $\sigma_{i,l}$ is a. best response for agent i.

b) For $i \in \mathbb{N}, j \in \mathbb{N}, i \neq j,\ if\ p_{i,y,j} = \psi_i(y,j), \omega = F(\sigma_1, \sigma_2, \dots, \sigma_j, \dots, \sigma_n)$, then $\Psi_i(\omega, \psi_i(y,j))$=$\psi_i(y,j)$.

That is, the outcome of the game is consistent with the prior belief system.

When a) and b) are satisfied, these strategies, belief systems, and outcomes constitute an equilibrium. The equilibrium is a stability condition, in that

1. If a) is violated, then i has not chosen a best response and can be expected to correct this error by choosing strategy $\sigma_{i,k}$ or still another with a preferred outcome.
2. If b) is violated, then i has made an error in his judgement of the type of player j and may be expected to revise his belief system, perhaps by applying Bayes' rule.

The stable equilibrium then corresponds to a stable system of strategic decisions *and stable belief systems.*

Remark: One might expect these formal assumptions to be followed by some theorems, but that is not the objective of the present paper. Rather, the objective is to establish existence of a class of equilibria with stable mistaken beliefs as consistent with these assumptions, and the examples in the foregoing sections are offered as instances of such a class. In the examples, some probability distributions are degenerate. Unbelievers' prior probabilities are assumed to be the correct ones and so observed events will be consistent with them, no revision being needed. For believers, the conditional probability that a labeled agent is of type A is zero, and Bayes' rule does not yield any revision of a zero conditional probability regardless of observed evidence. Nevertheless the observation of evidence inconsistent with a hypothesis is sufficient to establish that the correct probability is not one. Perhaps the best way to resolve this is to suppose that real agents resolve evidence in ways that are only approximately Bayesian.

iv. Some Lessons

Despite some relatively novel assumptions, these models are much simpler than systems of social stratification. Emotions and respect or the lack of it play no part in the models, which assume only rational decisions. There do seem to be some lessons to be drawn, however.

1) The beliefs of some agents about labelled agents may influence the rational decisions of the labelled agents in ways that are to the disadvantage of *both*.
2) The rational decisions of the labelled agents may be consistent with the false beliefs of unlabelled agents about them, stabilizing beliefs that are objectively false[7].
3) The false beliefs of some unlabelled agents may influence the rational decisions of other unlabelled agents (who do not hold the false beliefs) in ways that are to the disadvantage of all, and particularly of labelled agents.
4) The presence of unlabeled agents in the population who hold the false beliefs limits the effectiveness of both mixed and unmixed groupings of agents, but the limits to the effectiveness of mixed groupings are worse.
5) Given the possibility to choose, both unlabeled unbelievers and labeled agents may have reason to choose unmixed groupings. (Note Alsan, Garrick and Graziani 2019, Shari, Logan, and Miloucheva 2023).
6) "Statistical discrimination" is not innocent. By "statistical discrimination" we mean the construction of decision criteria using observations that exclude the label but that are correlated with the label. In the examples given above, this correlation is endogenous, and simply codifies the closed circle that stabilizes the false beliefs. (Compare Rose, p. 907 et seq..)

---

[7] But see Bohren, Imas, and Rosenberg (2019).

7) Quotas may work. If the quotas change the experiences of unlabeled agents in ways that break the stable circle of beliefs (and this is a big if!) then the quotas could improve the decisions made.

v. Concluding Summary

The concept of a stable noncooperative (Nash) equilibrium can be extended to cases in which agents differ in some details of the play of the game, such as payoffs or strategy sets. They are thought of as different “types,” and agents may be uncertain about the types of other players in the game. To this well established equilibrium theory the paper adds belief systems, by which a player may infer (correctly or not) the types of other players in the game. The concept of a stable belief system is formalized and a stable situation is one in which both the belief systems and the strategic decisions (as in Nash equilibrium) are stable. But stability does not imply truth. Examples have been given in which objectively false belief systems are stable. These examples involve the belief on the part of some agents that other agents who are distinguished by a visible or socially constructed signal are systematically more effort-averse than others, a belief that often seems to be associated with social stratification. The beliefs are stable although they are false.

Bibliography

Alsan, Marcella and Owen Garrick and Grant Graziani (2019), "Does Diversity Matter for Health? Experimental Evidence from Oakland ," The American Economic Review v. 109, no. 12 pp. 4071-4111.

Aumann, R. J. (1976), "Agreeing to Disagree," Annals of Statistics v. 4, pp. 1236-1239.

Aumann, R. J. (1999), "Interactive Epistemology I and II," International Journal of Game Theory v. 28, pp. 263-314.

Björklund , Anders and Markus Jäntti (1997), "Intergenerational Income Mobility in Sweden Compared to the United States ," The American Economic Review v. 87, no. 5 (Dec) pp. 1009-1018.

Blanchflower, David G and Phillip B Levine, and David J Zimmerman. (2003), "Discrimination in the Small-business Credit Market," Review of Economics and Statistics v. 85, no. 4 pp. 930-943.

Bohren, J. Aislinn and Alex Imas, and Michael Rosenberg (2019), "The Dynamics of Discrimination: Theory and Evidence," The American Economic Review v. 109, no. 10 pp. 3395-3436.

Bordalo, Pedro and Katherine Coffman, Nicola Gennaioli and Andrei Shleifer (2016), "STEREOTYPES ," The Quarterly Journal of Economics v. 131, no. 4 pp. 1753-1794.

Clark, Gregory et al. (2012), "Medieval England: Mobility in the Feudal Age," The Son Also Rises: Surnames and the History of Social Mobility (Princeton University Press).

Clark, Gregory and Neil Cummins (2014), "Surnames and Social Mobility in England, 1170-2012," Human Nature v. 25, pp. 517-537.

Clark, Gregory and Neil Cummins (2015), "INTERGENERATIONAL WEALTH MOBILITY IN ENGLAND, 1858-2012: SURNAMES AND SOCIAL MOBILITY," The Economic Journal v. 125, no. 582 (Feb) pp. 61-85.

Dorsinville, Jean L. (2023), "The Readers Write," (New York Times) available at: https://www.nytimes.com/2023/06/23/opinion/biden-industrial-policy.html, as of June 26, 2023, originally posted on 6/23/2023.

Garcia, Raffi E. and William A. Darity, Jr. (2022), "Self- Reporting Race in Small Business Loans: A Game-Theoretic Analysis of Evidence from PPP Loans in Durham, NC," AEA Papers and Proceedings v. 112, pp. 299-302.

Gibbons , Robert (1997), "An Introduction to Applicable Game Theory," The Journal of

Economic Perspectives v. 11, no. 1 (winter) pp. 127-149.

GUELL, MAIA and JOSE V. RODRIGUEZ MORA and CHRISTOPHER I. TELMER (2015), "The Informational Content of Surnames, the Evolution of Intergenerational Mobility, and Assortative Mating,"The Review of Economic Studies v. 82, no. 2 (April) pp. 693-735.

Harsanyi, John C. (1967), "Games with Incomplete Information Played by 'Bayesian' Players, Parts I, v. 14, pp.159-183; II, pp. 320-334; III, pp 486-502," Management Science v. 14

Harsanyi, John and Reinhard Selten (1988), A General Theory of Equilibrium Selection in Games (Cambridge: MIT Press).

Ichikawa, Jonathan Jenkins and Matthias Steup, "The Analysis of Knowledge", *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>.

Long, Jason and Joseph Ferrie (2013), "Intergenerational Occupational Mobility in Great Britain and the United States Since 1850," The American Economic Review v. 103, no. 4 pp. 1109-1137.

Lyonnet , Victor and Léa Stern (2023), "Stereotypes about Successful Entrepreneurs," AEA Papers and Proceedings v. 113, pp. 186-190.

McCain, Roger A. (2018), The Economics of Small Business: An Introductory Survey (Singapore, World Scientific).

McCain, Roger A. (2023), Game Theory: A Nontechnical Introduction to the Analysis of Strategy, Fourth Edition (World Scientific).

Nunn, Nathan (2022), "On the Dynamics of Human Behavior: The Past, Present, and Future of Culture, Conflict, and Cooperation," AEA Papers and Proceedings v. 112, no. 37 pp. 15-37).

Rose, Evan K. (2023), "A Constructivist Perspective on Empirical Discrimination Research," Journal of Economic Literature v. 61, no. 3 pp. 906-923.

Sen, Amartya K (1997), "From Income Inequality to Economic Inequality," Southern Economic Journal v. 64, no. 2 (Oct.) pp. 383-401.

Shari, J. Eli and Trevon D. Logan, and Boriana Miloucheva (2023), "The Enduring Effects of Racial Discrimination on Income and Health," Journal of Economic Literature v. 61, no. 3 pp. 924-940.

TORCHE, FLORENCIA (2015), "Analyses of Intergenerational Mobility: An Interdisciplinary Review," The Annals of the American Academy of Political and Social Science v. 657, (January) pp. 37-62.