DEF 25, 33 – 45 (2002)

Decisions in Economics and Finance

Bargaining over an uncertain outcome: the role of beliefs

Antoine Billot[1], **Alain Chateauneuf**[2], **Itzhak Gilboa**[3], **Jean-Marc Tallon**[4]

[1] CERAS-ENPC, Université Paris II and Institut Universitaire de France
e-mail: billot@u-paris2.fr
[2] CERMSEM, Université Paris I Panthéon-Sorbonne
e-mail: chateaun@univ-paris1.fr
[3] Eitan Berglas School of Economics and Recanati School of Business
e-mail: igilboa@post.tau.ac.il
[4] EUREQua, CNRS, Université Paris I Panthéon-Sorbonne and Università di Venezia
e-mail: jmtallon@univ-paris1.fr

Received: 10 May 2001 / Accepted: 22 August 2001

Abstract. We study the Nash bargaining solution of a problem in which two agents bargain over an uncertain outcome. Under the assumptions of risk neutrality and of constant absolute risk aversion, we study the way that the solution varies, ex ante, when we vary the beliefs of one agent. Changing an agent's beliefs in a way that makes them "more distant" from the other agent's beliefs makes the second agent better off.

Mathematics Subject Classification (2000): 91A12, 91B26

Journal of Economic Literature Classification: C71, C78

1. Introduction

Economic theory typically assumes that agents share a prior probability over states of the world. This "common prior assumption" (CPA) is theoretically appealing for several reasons (see Aumann (1987)). In particular, when the CPA fails to hold, some agents entertain "incorrect" beliefs in the eyes of others, and, should the model specify a prior probability as well, also in the eyes of the modeler. Yet casual observation indicates that people differ in their beliefs. Moreover, the CPA, coupled with certain assumptions

We thank two referees for helpful comments.

of rationality and common knowledge thereof, yields counter-intuitive results that preclude agreeing to disagree, betting and trading. (See Aumann (1976), Geanakoplos and Polemarchakis (1982), and others. See Geanakoplos (1994) for a survey.[1]) It is therefore of interest to study markets in which different agents might have different beliefs (see Morris (1995)).

Different beliefs induce trade. Just as agents with different tastes may benefit from trade, so may agents with different beliefs, where trade occurs ex ante. Intuitively, agents whose tastes are "more different" from each other have more to gain by trading. By a similar token, agents whose beliefs are "more different" from each other stand to gain more by trading with each other. This paper attempts to formalize this intuition using a cooperative approach. Such an approach, which does not specify the actual non-cooperative game that is being played, promises to provide some limited yet robust insights. Specifically, we adopt the Nash bargaining solution (Nash (1953)) and study how it varies as a function of beliefs.[2] We study the simplest case of two agents who bargain over an uncertain financial asset; that is, an allocation specifies the amounts of a single commodity, say, money, that each of the two agents obtains in each state of the world. The agents are assumed to have vNM utility functions exhibiting either risk neutrality or constant absolute risk aversion.[3] Under these assumptions, we study the way in which the Nash bargaining solution (NBS) varies, ex ante, when we vary the beliefs of one agent. Specifically, we ask whether the other agent, whose beliefs have not changed and can therefore be used for comparison, is better or worse off at the NBS after the change.

The NBS for two risk neutral agents turns out to be rather different from the solution for the case of two risk averse agents (with constant absolute risk aversion). Yet, in both cases, we find a similar result: changing agent 1's beliefs in a way that makes them more "distant" from the beliefs of agent 2 makes agent 2 better off at the NBS. The definition of "more distant" beliefs is somewhat more cumbersome in the case of risk averse agents. But, in both cases, the basic intuition is supported by our results: it is easier to bargain with someone who has different beliefs simply because different beliefs generate a larger pie.

The rest of this paper is organized as follows. In Section 2 we present the simple set-up. Section 3 deals with the risk neutral case, whereas Section 4

[1] Admittedly, some of these counter-intuitive results, such as Milgrom and Stokey (1982), do not depend on the CPA.

[2] Needless to say, one may study other cooperative solution concepts, as well as non-cooperative models. The Nash bargaining solution, arguably the most popular cooperative concept, constitutes a reasonable starting point for studying the role of beliefs in bargaining theory.

[3] In this paper, we consider only cases in which both agents are risk neutral or both are risk averse.

deals with the case of risk aversion, with constant absolute risk aversion. Section 5 concludes.

2. The set-up

We consider a bargaining problem in which two agents have to split a random pie. Let $S = [0, 1]$ be the state space and $w : S \to \mathbb{R}_{++}$ be a measurable bounded function representing the pie to be split, i.e., $w(s)$ is the size of the pie in state s. The two agents have (possibly) different beliefs over S. Let $p(.)$ and $q(.)$ be densities over S representing agent 1's and 2's beliefs, respectively. We assume throughout that w, p and q are continuous and strictly positive.

Let $x : [0, 1] \to \mathbb{R}_+$ be a measurable function satisfying $x(s) \leq w(s)$, interpreted as follows: $x(s)$ is what agent 1 obtains in state s, and $w(s)-x(s)$ is what agent 2 obtains in state s. The expected utility of a given split is given by:

$$u_1(x) = \int_0^1 p(s)v_1(x(s))ds$$

$$u_2(x) = \int_0^1 q(s)v_2(w(s) - x(s))ds$$

where $v_i : \mathbb{R}_+ \to \mathbb{R}$ is agent i's von Neumann–Morgenstern utility index.

We assume in the following that each v_i exhibits constant absolute risk aversion. Hence it can take one of two forms: it may be linear, $v_i(x) = x$, or it may be exponential, i.e., $v_i(x) = -\exp(-a_i x)/a_i$ with $a_i > 0$. The solution of the Nash bargaining problem turns out to take a rather different form in the two cases.

3. The risk neutral case

We first focus on the case of risk neutrality: $v_1(x) = x$ and $v_2(w - x) = w - x$.

The frontier of the set of feasible utilities (i.e., the set of Pareto optimal utility outcomes) can be described by a collection of optimization problems. One such collection is given by the following family, parametrized by K:

$$\max_x \int_0^1 p(s)x(s)ds$$

$$\text{s.t.} \begin{cases} \int_0^1 q(s)(w(s) - x(s))ds = K \\ 0 \leq x(s) \leq w(s) \ \ \forall s, \end{cases}$$

Alternatively, the frontier can be described by varying $\gamma > 0$ in the following program (see Lehrer and Pauzner (1999) who employ the same description):

$$\max_{\{x|0\le x(s)\le w(s)\ \forall s\}} \int_0^1 p(s)x(s)ds + \gamma \int_0^1 q(s)(w(s) - x(s))ds,$$

that is,

$$\max_{\{x|0\le x(s)\le w(s)\ \forall s\}} \int_0^1 (p(s) - \gamma q(s))x(s)ds.$$

The solution to this problem is trivially:

$$\begin{cases} x^\star(s) = 0 & \text{if } p(s) < \gamma q(s) \\ x^\star(s) \in [0, w(s)] & \text{if } p(s) = \gamma q(s) \\ x^\star(s) = w(s) & \text{if } p(s) > \gamma q(s) \end{cases}.$$

Now define a pre-order on S as follows[4]: $s \succ s'$ if and only if $p(s)/q(s) > p(s')/q(s')$, or, equivalently, $p(s)/p(s') > q(s)/q(s')$. Assume that $s \succ s'$ implies $s > s'$, that is, that p/q is non-decreasing over $[0, 1]$. This assumption simplifies notation without any loss of generality.

Define $U_1(\alpha) = \int_\alpha^1 p(s)w(s)ds$, that is, agent 1's expected utility if she gets all the pie in all states from α to 1 and zero in the other states, and $U_2(\alpha) = \int_0^\alpha q(s)w(s)ds$, that is, agent 2's expected utility of getting the entire pie in states 0 to α and nothing otherwise. Since p, q and w are strictly positive, U_1 is decreasing and U_2 is increasing. The Pareto frontier (in the space of utility) is given by[5]

$$\{(U_1(\alpha), U_2(\alpha))|\alpha \in [0, 1]\}.$$

Intuitively, in the presence of risk neutrality, Pareto optimality dictates that, in each given state, one agent gets the entire pie. The cutoff state at which we switch from giving everything to agent 2 to giving everything to agent 1 is given by the weight of each agent in the social welfare function. Observe that the optimal allocation is not an interior solution (state by state).

---

[4] An alternative approach to the problem, suggested to us by the editor, would be to assume that w depends on a random variable σ taking values in $[0, 1]$. Players may then have different state spaces S_1 and S_2. To player 1 (resp., 2) σ is a function from S_1 (resp., S_2) to $[0, 1]$ which we will denote by σ_1 (resp. σ_2). Players then "contract" on a pie whose size depends on the observable σ. Consider now the following stochastic orders: let σ_1 and σ_2 have c.d.f. F and G respectively with densities f and g and common support. Then g/f increases over the support if and only if $G \circ F^{-1}$ is convex (see Shaked and Shantikumar (1994)).

[5] If p/q is not increasing the explicit expressions for u_1 and u_2 would have to change so that the domain of integration for u_1 is $\{s|p(s)/q(s) \ge \alpha\}$ and the complement for u_2.

The domain of Pareto optimal allocations thus defined is convex. Indeed, u_i, the expected utility functional of player i, is linear in x for $i = 1, 2$. Therefore one can find the NBS for this problem by solving the following problem, where we assume that the disagreement point is zero in every state (i.e., if the bargaining were to fail, the two agents would receive a payoff of zero in every state):

$$\max_{\alpha \in [0,1]} \left(\int_\alpha^1 p(s)w(s)ds \right) \times \left(\int_0^\alpha q(s)w(s)ds \right).$$

The first order condition of this problem is

$$-p(\alpha)w(\alpha)U_2(\alpha) + q(\alpha)w(\alpha)U_1(\alpha) = 0$$

i.e., the NBS $\alpha^\star$ is given by

$$\frac{U_1(\alpha^\star)}{U_2(\alpha^\star)} = \frac{p(\alpha^\star)}{q(\alpha^\star)}.$$

The problem now is to characterize how this solution is affected by changes in the agents' beliefs. Specifically, we look at the following question: suppose that agent 1 has new beliefs given by $p' \equiv p + f$, where f has to satisfy $\int_0^1 f(s)ds = 0$. How does this affect agent 2's utility in the (new) NBS ?

Define a new bargaining problem where agent 2 is as above. Call $\widetilde{\alpha}$ the NBS of this new problem. Observe that, if $(p + f)/q$ is increasing in s (recall that p/q is non-decreasing in s) it is characterized by:

$$\frac{\widetilde{U}_1(\widetilde{\alpha})}{U_2(\widetilde{\alpha})} = \frac{(p+f)(\widetilde{\alpha})}{q(\widetilde{\alpha})},$$

where

$$\widetilde{U}_1(\widetilde{\alpha}) = \int_{\widetilde{\alpha}}^1 (p(s) + f(s))w(s)ds$$

Our first result is that, under risk neutrality, changes in agent 1's beliefs of this type increase agent 2's welfare.

Proposition 1. *Assume that f/p is non-decreasing with respect to p/q and that $p + f$ is continuous and strictly positive. Then $U_2(\widetilde{\alpha}) \geq U_2(\alpha^\star)$.*

The type of exercise we conduct here is reminiscent of the question of, under conditions of certainty, how agent 2's welfare is affected by a change in agent 1's utility function, or, more precisely, in the concavity of agent 1's utility function (Kihlstrom, Roth and Schmeidler (1981), Roth and Rothblum (1982)). Moreover, one could use their result to prove ours by showing that, with beliefs p', agent 1's new utility function, $\widetilde{u}_1$, is a concave transform of her utility with beliefs p, u_1. Here we prove our result directly.

Proof. First note that under our notation assumption p/q is non-decreasing in s. Further, since f/p is non-decreasing in p/q, then $(p+f)/q$ is also non-decreasing in s. Hence, $\widetilde{\alpha}$ is given by

$$\frac{\widetilde{U}_1(\widetilde{\alpha})}{U_2(\widetilde{\alpha})} = \frac{(p+f)(\widetilde{\alpha})}{q(\widetilde{\alpha})}.$$

Moreover, $f(s)/p(s) \geq f(\alpha^\star)/p(\alpha^\star)$ for all $s \in [\alpha^\star, 1]$; hence

$$f(s)w(s) \geq \frac{f(\alpha^\star)}{p(\alpha^\star)} p(s)w(s) \ \ \forall s \in [\alpha^\star, 1],$$

which implies that

$$\int_{\alpha^\star}^1 f(s)w(s)ds \geq \frac{f(\alpha^\star)}{p(\alpha^\star)} \int_{\alpha^\star}^1 p(s)w(s)ds.$$

Hence, $\int_{\alpha^\star}^1 (p(s)+f(s))w(s)ds \geq \int_{\alpha^\star}^1 p(s)w(s)ds + \frac{f(\alpha^\star)}{p(\alpha^\star)} \int_{\alpha^\star}^1 p(s)w(s)ds$ and then

$$\frac{p(\alpha^\star)}{q(\alpha^\star)} \frac{\int_{\alpha^\star}^1 (p(s)+f(s))w(s)ds}{\int_{\alpha^\star}^1 p(s)w(s)ds} \geq \frac{(p+f)(\alpha^\star)}{q(\alpha^\star)}.$$

Now observe that $\int_{\alpha^\star}^1 (p(s)+f(s))w(s)ds = \widetilde{U}_1(\alpha^\star)$ and that $\int_{\alpha^\star}^1 p(s)w(s)ds = U_1(\alpha^\star) = \frac{p(\alpha^\star)}{q(\alpha^\star)} U_2(\alpha^\star)$. Hence,

$$\frac{\widetilde{U}_1(\alpha^\star)}{U_2(\alpha^\star)} \geq \frac{(p+f)(\alpha^\star)}{q(\alpha^\star)}.$$

Now, $\frac{\widetilde{U}_1(s)}{U_2(s)}$ is a decreasing function while $\frac{(p+f)(s)}{q(s)}$ is non-decreasing in s. Since $\widetilde{\alpha}$ is defined by

$$\frac{\widetilde{U}_1(\widetilde{\alpha})}{U_2(\widetilde{\alpha})} = \frac{(p+f)(\widetilde{\alpha})}{q(\widetilde{\alpha})},$$

we see that

$$\frac{\widetilde{U}_1(\alpha^\star)}{U_2(\alpha^\star)} \geq \frac{(p+f)(\alpha^\star)}{q(\alpha^\star)} \Rightarrow \widetilde{\alpha} \geq \alpha^\star \Rightarrow \begin{cases} U_1(\widetilde{\alpha}) \leq U_1(\alpha^\star) \\ U_2(\widetilde{\alpha}) \geq U_2(\alpha^\star) \end{cases}.$$

Hence, changing agent 1's beliefs from p to $p+f$ cannot decrease agent 2's welfare whenever f/p is non-decreasing. □

Observe that, since f/p is non-decreasing, adding f to p can be viewed as "pushing" p further away from q:

$$\frac{(p+f)(s)}{(p+f)(s')} \geq \frac{p(s)}{p(s')} \geq \frac{q(s)}{q(s')}.$$

Hence, Proposition 1 establishes that, whenever agent 1's beliefs become "more distant" from agent 2's beliefs, then, at the NBS, agent 2's welfare increases if agents are risk neutral.

How does agent 2 take advantage of such changes in agent 1's beliefs? One possible story runs as follows. The states that agent 1 considered more likely than did agent 2 are now considered even more likely by her. Hence, agent 1 puts more weight on states that agent 2 considers as relatively less likely. Giving the pie to agent 1 in those states increases her utility relatively more than before and allows agent 2 to obtain the pie in more states than before.

Example 1. Let $w(s) = s$, and $p(s) = q(s) = 1 \ \forall s \in [0, 1]$, and observe that the natural order on $[0, 1]$ is consistent with the preorder induced by p/q. The Nash bargaining solution is given by $\alpha^\star = \sqrt{2}/2$, the solution to the problem

$$\max_{\alpha} \int_{\alpha}^{1} s ds \int_{0}^{\alpha} s ds.$$

Now suppose that $p' = p + f$ where $f(s) = s - (1/2)$, that is, $p'(s) = s + (1/2)$. It is easy to see that f/p is increasing in s. Hence, the NBS is given by the solution to the problem

$$\max_{\alpha} \int_{\alpha}^{1} s(s + (1/2)) ds \int_{0}^{\alpha} s ds.$$

The first order condition yields

$$\left[-\alpha^2 - \frac{\alpha}{2}\right]\frac{\alpha^2}{2} + \left[\frac{7}{12} - \frac{\alpha^3}{3} - \frac{\alpha^2}{4}\right]\alpha = 0,$$

that is,

$$g(\alpha) \equiv \frac{5}{3}\alpha^3 + \alpha^2 - \frac{7}{6} = 0.$$

It is straightforward to check that g is increasing on $[0, 1]$ and that $g(\sqrt{2}/2) < 0$, while $g(1) > 0$. Hence, the NBS of the problem with modified beliefs p' is greater than $\sqrt{2}/2$ and agent 2 (whose beliefs were not modified) is better off at this new solution.

4. The Nash bargaining solution with CARA utility functions

We now turn to the case in which both agents are risk averse and have constant absolute risk aversion. Specifically, let

$$v_1(x) = \frac{-e^{-ax}}{a} \quad \text{and} \quad v_2(x) = \frac{-e^{-bx}}{b}$$

with $a > 0$ and $b > 0$.

We start by characterizing the possibility domain in the space of utility allocations. This time we use the following collection of optimization problems:

$$\begin{aligned} \max_x & \int_0^1 \frac{-e^{-ax(s)}}{a} p(s)ds \\ \text{s.t.} & \int_0^1 \frac{-e^{-b(w(s)-x(s))}}{b} q(s)ds = u_2 \end{aligned}$$

for any $u_2 \leq 0$.

With λ denoting the Lagrange multiplier associated to the constraint, the first order conditions for this problem yield:

$$\begin{aligned} \frac{p(s)}{q(s)} \frac{e^{-ax(s)}}{e^{-b(w(s)-x(s))}} = \lambda \Rightarrow e^{-(a+b)x(s)} &= \lambda \frac{q(s)}{p(s)} e^{-bw(s)} \\ &\Rightarrow \begin{cases} x(s) = \frac{1}{a+b} \log \frac{p(s)}{q(s)} - \frac{\log \lambda - bw(s)}{a+b} \\ w(s) - x(s) = \frac{1}{a+b} \log \frac{q(s)}{p(s)} + \frac{\log \lambda + aw(s)}{a+b}. \end{cases} \end{aligned}$$

Hence,

$$\begin{aligned} u_1(x) &= \int_0^1 \frac{-e^{-ax(s)}}{a} p(s)ds \\ &= \frac{-1}{a} \int_0^1 e^{\frac{-a}{a+b} \log \frac{p(s)}{q(s)} + a \frac{\log \lambda - bw(s)}{a+b}} p(s)ds \\ &= \frac{-1}{a} \int_0^1 \lambda^{\frac{a}{a+b}} e^{\frac{-ab}{a+b} w(s)} p(s)^{\frac{b}{a+b}} q(s)^{\frac{a}{a+b}} ds. \end{aligned}$$

Similarly,

$$\begin{aligned} u_2(x) &= \int_0^1 \frac{-e^{-b(w(s)-x(s))}}{b} q(s)ds \\ &= \frac{-1}{b} \int_0^1 \lambda^{-\frac{b}{a+b}} e^{-\frac{ab}{a+b} w(s)} p(s)^{\frac{b}{a+b}} q(s)^{\frac{a}{a+b}} ds. \end{aligned}$$

Let $K = \int_0^1 e^{-\frac{ab}{a+b}w(s)} p(s)^{\frac{b}{a+b}} q(s)^{\frac{a}{a+b}} ds$ and observe that, at a solution,

$$u_1(x^\star) = u_1 = -\frac{\lambda^{\overline{a+b}}{}^{a}}{a} K$$

$$u_2(x^\star) = u_2 = -\frac{\lambda^{-\frac{b}{a+b}}}{b} K,$$

and, therefore, the Pareto frontier in the (u_1, u_2) plane is given by

$$(-u_1) = \frac{K}{a}\left(\frac{K}{b}\right)^{\frac{a}{b}} (-u_2)^{-\frac{a}{b}}.$$

We now turn to the Nash bargaining solution of this problem, taking the disagreement point to be 0 in all states. Thus, we are led to maximize $(u_1-(-1/a))\times(u_2-(-1/b))$. Knowing the relationship, at a Pareto optimal solution, between u_1 and u_2, we can replace u_1 by $\frac{-K}{a}\left(\frac{K}{b}\right)^{\frac{a}{b}}(-u_2)^{\frac{-a}{b}}$ and solve the following problem for u_2:

$$\max_{u_2} \left(-\frac{K}{a}\left(\frac{K}{b}\right)^{\frac{a}{b}} (-u_2)^{-\frac{a}{b}} + \frac{1}{a}\right) \times \left(-(-u_2) + \frac{1}{b}\right).$$

Taking the first order condition leads, after some computation, to the following implicit equation for $(-u_2)$:

$$\frac{K(p,q)}{a}\left(\frac{K(p,q)}{b}\right)^{\frac{a}{b}} (-u_2)^{-\frac{a}{b}} \left[\frac{a}{b^2}(-u_2)^{-1} + \frac{b-a}{b}\right] = \frac{1}{a},$$

where we write $K(p,q)$ to emphasize that K depends on both agents' beliefs.

Note that the above expression is positive and increasing in $K(p,q)$. Furthermore, observe that, on the domain where it is positive, $h(u_2) \equiv (-u_2)^{-\frac{a}{b}}\left[\frac{a}{b^2}(-u_2)^{-1} + \frac{b-a}{b}\right]$ is an increasing function. Hence, any change in agent 1's beliefs (p) that decreases $K(p,q)$ leads to an increase in u_2.

Define the following preorder on S. Say that $s' \succ s$ if

$$\frac{p(s')}{q(s')} > \frac{p(s)}{q(s)} \frac{e^{-bw(s')}}{e^{-bw(s)}}.$$

Note that this preorder may not be consistent with the one used under risk neutrality. Further, observe that $s' \succ s$ iff, for the allocation $x(s) = 0$ for all s,

$$\frac{p(s')v_1'(x(s'))}{q(s')v_2'(w(s') - x(s'))} > \frac{p(s)v_1'(x(s))}{q(s)v_2'(w(s) - x(s))}$$

or

$$\frac{q(s)v_2'(w(s)-x(s))}{q(s')v_2'(w(s')-x(s'))} > \frac{p(s)v_1'(x(s))}{p(s')v_1'(x(s'))};$$

that is, $s' \succ s$ iff, if we start at an allocation where agent 2 gets the entire pie, agent 1's marginal rate of substitution between income at state s and at state s' is larger than that of agent 2. Note that the allocation $x(s) = 0$ for all s defines a point which is on the Pareto frontier both before and after the change in agent 1's beliefs.

We now prove that any modification of p that transfers probability mass from states that are ranked relatively low by the $\succ$ scale to states that are ranked relatively high by the $\succ$ scale increases agent 2's utility.

Proposition 2. *Let p' be such that $p'(s) = p(s) + f(s)$ where $\int_0^1 f(s)ds = 0$ and f is monotone increasing with respect to $\succ$ and assume that p' is continuous and strictly positive. Then, agent 2 is better off at the Nash bargaining solution when agent 1 has beliefs p' than when she has beliefs p.*

Proof. To prove the proposition it will suffice to consider local changes of the following type. Let $s_2 \succ s_1$. Observe that this implies that $s_2 \neq s_1$. Since $\frac{p(s)}{q(s)}\frac{1}{e^{-bw(s)}}$ is continuous, we can find a $\delta > 0$ such that, for every t_2 in a δ-neighborhood of s_2 and every t_1 in a δ-neighborhood of s_1,

$$\frac{p(t_2)}{q(t_2)} > \frac{p(t_1)}{q(t_1)}\frac{e^{-bw(t_2)}}{e^{-bw(t_1)}}.$$

Let p' be such that: (i) $p'(s) = p(s) - \varepsilon$ in the δ-neighborhood of s_1; (ii) $p'(s) = p(s) + \varepsilon$ in the δ-neighborhood of s_2; and (iii) $p'(s) = p(s)$ outside these neighborhoods. We will show that with this definition, agent 2 is better off at the NBS when agent 1 has beliefs p' than he is when agent 1 has beliefs p.

We wish to show that K is a decreasing function of ε. Approximating p, q and w on the neighborhoods of s_1 and of s_2 by their respective values at s_1 and at s_2 (respectively), one obtains

$$\frac{dK(p,q)}{d\varepsilon} < 0 \Leftrightarrow \left(\frac{q(s_1)}{p(s_1)}\right)^{\frac{a}{a+b}} e^{-\frac{ab}{a+b}w(s_1)} > \left(\frac{q(s_2)}{p(s_2)}\right)^{\frac{a}{a+b}} e^{-\frac{ab}{a+b}w(s_2)}$$
$$\Leftrightarrow \frac{p(s_1)}{q(s_1)} < \frac{p(s_2)}{q(s_2)}\frac{e^{-bw(s_1)}}{e^{-bw(s_2)}}.$$

Hence, if $s_2 \succ s_1$, then $\frac{dK(p,q)}{d\varepsilon} < 0$ and hence agent 2's utility increases.

It is straightforward to show that a density p' such that $p'(s) = p(s) + f(s)$ can be approximated by a sequence of ε-transfers of the type discussed above. □

Remark 1. Contrary to the risk neutral case, the solution here is interior in each state, i.e., each agent gets a positive fraction of the pie in each state.

Remark 2. If the agents start out with identical beliefs, then the order on s with respect to which f has to be monotone (in the risk averse case) depends only on w and not on the beliefs of the agents. Explicitly, if $p = q$, we have $s' \succ s$ iff $w(s') > w(s)$. That is, in case of identical initial beliefs, agent 2 will be better off at the NBS if agent 1 puts relatively more probability on the favorable states of the world.

Observe that, when the two agents have the same utility function (i.e., the same constant absolute risk aversion coefficient, that is, $a = b$), the expression for the NBS takes the following simple form:

$$u_1 = u_2 = -\frac{1}{a}\int_0^1 e^{-a\frac{w(s)}{2}}(p(s)q(s))^{1/2}ds,$$

where $e^{-a\frac{w(s)}{2}}$ is the utility of consuming half of the pie in state s and $(p(s)q(s))^{1/2}$ is the geometric mean of the two agents' beliefs.

In that particular case, $x(s) = \dfrac{1}{2a}\log\dfrac{p(s)}{q(s)} + \dfrac{w(s)}{2}$ and $w(s) - x(s) = \dfrac{1}{2a}\log\dfrac{q(s)}{p(s)} + \dfrac{w(s)}{2}$. It is interesting to note that the absolute amount of the pie that goes to agent 2 in state s increases with $q(s)/p(s)$, that is, the same order as the one defined in the linear, risk neutral case. However, this does not translate to utility terms.

Example 2. Take the set-up of Example 1 and let $w(s) = s$ and $p(s) = q(s) = 1\ \forall s \in [0, 1]$. The order defined here is again the same as the natural order on [0, 1]. Indeed, $s' \succ s$ if

$$\frac{p(s')}{q(s')} > \frac{p(s)}{q(s)}\frac{e^{-bw(s')}}{e^{-bw(s)}},$$

that is, if $e^{-bw(s')} < e^{-bw(s)}$ and $w(s') > w(s)$.

If $a = b$, i.e., the agents have the same utility functions, the NBS is given by splitting the pie exactly in half in each state ($x(s) = w(s)/2$).

A change in agent 1's beliefs to p' where p' is as in the proof of Proposition 2, namely: (i) $p'(s) = p(s) - \varepsilon$ in a δ-neighborhood of s; (ii) $p'(s) = p(s) + \varepsilon$ in a δ-neighborhood of s'; and (iii) $p'(s) = p(s)$ outside these neighborhoods, where δ is smaller than $\frac{s'-s}{2}$, increases agent 2's welfare.

5. Discussion

That different tastes give rise to gains from trade is accepted as a fundamental economic insight. But this is not the case when it comes to different beliefs. After all, when two agents trade only thanks to differing beliefs, it is as if they were betting against each other. Must not one of them be wrong about her beliefs? Is it not the case that one of the agents is fooling the other? Alternatively, will the agents not learn from experience and eventually converge to the same, true beliefs? Moreover, isn't the very fact that one wishes to trade a signal that this trade should not be carried out?

Each of these questions, if answered in the affirmative, may render our exercise useless. Indeed, we find that in many contexts some of these questions might be so answered. But not always. We first consider the issue of "true" or "objective" beliefs. There are situations where these can be well-defined, as in beliefs about a roulette wheel. In these situations our analysis can only be relevant if at least one agent is simply wrong, in that she ascribes to events subjective probabilities that differ from available or computable objective probabilities of the same events. But objective probabilities seem to be the exception rather than the rule. Most macro-economically relevant events are too novel to be assigned objective probabilities. For most of the events that affect stock market trading one cannot ascribe probabilities in any "scientific" or objective way.

Subjective probabilities may converge by Bayesian update. This is true if the same situation repeats itself in the same manner over and over again, such as with the toss of a coin. Having a sufficiently diffused prior probability on the parameter of the coin (say, the probability of "Head"), a Bayesian agent will entertain beliefs that eventually converge to the true parameter. But, in many economically relevant situations, the source of uncertainty is not encountered enough times for convergence to occur. Again, most macro-economic phenomena are unique. By a similar token, events that are relevant to stock market trading never repeat themselves in exactly the same way. Convergence of Bayesian beliefs is therefore possible only under rather restrictive assumptions, say, that the prior beliefs over a very large space are absolutely continuous with respect to each other (see Blackwell and Dubins (1962), Kalai and Lehrer (1993)).

Still, one may argue, even without convergence, rational agents should not trade. If the agents incorporate other agents' willingness to trade into their state space, Bayesian updating should lead to no-trade results at the very first stage. Intuitively, the very fact that one is willing to trade says that she knows something we do not, and that, consequently, we should not be trading with her. Our analysis completely ignores this strategic aspect. Indeed, agents' actions and their willingness to trade are not part of the description of the state of the world in our model. In that, our agents exhibit bounded rationality.

By contrast, models that do incorporate these strategic aspects, such as that of Aumann (1987), assume a very high degree of rationality and of common knowledge thereof. The volume of stock market trade seems to suggest that these assumptions are too extreme. Future research might provide models that will incorporate some degree of strategic reasoning, without assuming common knowledge of rationality.

References

Aumann, R. (1976): Agreeing to disagree. *The Annals of Statistics* **4**, 1236–1239

Aumann, R. (1987): Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* **55**, 1–18

Blackwell, D., Dubins, L. (1962): Merging of opinions with increasing information. *The Annals of Mathematical Statistics* **33**, 882–886

Geanakoplos, J. (1994): Common knowledge. In: Aumann, R., Hart, S. (eds.): *Handbook of game theory with economic applications*. Vol. II, North-Holland, Amsterdam, pp. 1437–1496

Geanakoplos, J., Polemarchakis, H. (1982): We can't disagree forever. *Journal of Economic Theory* **28**, 192–200

Kalai, E., Lehrer, E. (1993): Rational learning leads to Nash equilibrium. *Econometrica* **61**, 1019–1045

Kihlstrom, R., Roth, A., Schmeidler, D. (1981): Risk aversion and solutions to Nash's bargaining problem. In: Moeschlin, O., Pallaschke, D. (eds.): Game theory and mathematical economics. North-Holland, Amsterdam, pp. 65–71

Lehrer, E., Pauzner, A. (1999): Repeated games with differential time preferences. *Econometrica* **67**, 393–412

Milgrom, P., Stokey, N. (1982): Information, trade and common knowledge. *Journal of Economic Theory* **26**, 17–27

Morris, S. (1995): The common prior assumption in economic theory. *Economics and Philosophy* **11**, 227–253

Nash, J. (1953): Two-person cooperative games. *Econometrica* **21**, 128–140

Roth, A., Rothblum, U. (1982): Risk aversion and Nash's solution for bargaining games with risky outcomes. *Econometrica* **50**, 639–647

Shaked, M., Shantikumar, J.G. (1994): Stochastic orders and their applications. Academic Press, Boston