

CARESS Working Paper 98-03

Trade Between Rational Agents as a Result of Asymmetric Information

Yoram Halevy*

University of Pennsylvania

Department of Economics

Philadelphia PA 19104

April 1998

Abstract

“No Trade” Theorems claim that the mere arrival of new information can not induce trade between rational agents, even in the presence of asymmetric information. We analyze an economy in which the information agents receive is without noise. As long as preferences abide by the Sure Thing Principle, the no-trade result holds. However, under more general preferences, which do not satisfy the Sure Thing Principle, we find that dynamic consistency is a necessary but might not be a sufficient condition for information not to induce trade, even with noiseless information. We provide sufficient conditions for a No-Trade

Theorem and investigate the pattern of trade with and without dynamic consistency of preferences.

JEL classification: D50, D81, D82

*This work is based on the first chapter in the author’s Ph.D. dissertation written under the supervision of Prof. Menahem Yaari at The Hebrew University of Jerusalem. I would like to thank him for his advice and support. I am grateful to Sergiu Hart and Oscar Volij for very helpful discussions. I benefitted from remarks and discussions on the subject with Luca Anderlini, Robert Aumann, Tomer Blumin, Yossi Feinberg, George Mailath and Asher Wolinsky. I thank Meggie Eisenstaedt for her help in editing this work.

Keywords: No-Trade Theorem, Sure Thing Principle, Common knowledge, Dynamic consistency, Consequentialism, Uncertainty aversion.

1. Introduction

Rationality, which is an underlying assumption of economic modeling, implies that people will trade only if it is beneficial to all participants. It is evident that a high volume of trade is driven by informational incentives (sometimes referred to as a “speculative trade”). Empirical works that support this observation are numerous (e.g. Kandel and Pearson, 1995; Lang *et al.*, 1992), but the theoretical explanation of why trade is driven by asymmetric information is indirect. Milgrom and Stokey (1982) showed that in an exchange economy

1. If the ex-ante allocation is Pareto-efficient, i.e., if the agents have exhausted all ex-ante benefits from trade; and it
2. Agents share a common prior belief on signals’ likelihood contingent on the contractible states; and if
3. Rationality is common knowledge, then

agents cannot agree on a trade that will dominate the ex-ante allocation. This result is a cornerstone in explaining trade, since it contrast with the empirical evidence, and forces us to validate each one of the assumptions underling it.

The immediate explanation of this apparent contradiction, is that the ex-ante allocation is not Pareto-efficient. This explanation is not valid in many real markets (as in the stock of foreign exchange markets), where the allocation mechanism is efficient, but most of the trade is speculative.

The second explanation, questions the common prior assumption inherent in the theorem. Although it is true that when priors differ trade could be achieved, Morris (1994) has shown that not just any difference suffices to induce trade. Moreover, assuming different priors calls for a deeper philosophical explanation. According to the “Harsanyi Doctrine” (Harsanyi, 1967/8), two agents who have access to the same information and the same training, ought to arrive at identical conclusions. In the case of probabilities, differences in the likelihood assessment of a certain event must be explained (according to this doctrine) by different information, since all rational agents use

Bayes' rule. On the other hand, one could ask why allow heterogeneity in tastes while we impose unanimity in prior beliefs. In this work we do not pursue this line of research.¹ Moreover, we assume that information is noiseless - and in this context, any heterogeneity in prior beliefs does not lead to trade.

The third postulate in the No-Trade Theorem is the *common knowledge of rationality*. Rationality is formally defined as maximizing an agent's welfare given available information. Although irrational behavior is commonplace, the question is how can we relax this assumption in economic modeling, and what are the building blocks of the common knowledge of rationality. One possible consistent explanation is that although every agent is rational, each once doubts the rationality of other agents with some positive (possibly small) probability. Neeman (1996) has shown that such doubts are sufficient to establish trade. Another explanation could rely on some flaw in the agents' information processing, e.g., the information sets are not a partition of the states of the world (e.g. Rubinstein and Wolinsky, 1990).

In this paper we take a different view of the rationality of the single agent. We try to bridge over the gap between the single-agent decision theoretic literature, and the economic literature which deals with the existence of trade as an informational phenomena. The objective of the paper is twofold: First, to view the existing result in a new, and simple, perspective; and second, to analyze accurately the economic consequences of departing from the standard assumptions regarding the agents' decision process. The latter analysis enable us to shed new light on the single decision theoretic literature.

We show that when information is noiseless (hence markets are complete) a sufficient condition for the No-Trade Theorem is Savage's *Sure Thing Principle*, which states that the preference between two contracts ("acts" in Savage's one-decision-maker terminology), depends solely on their consequences in states of the world in which they differ. In other words, preferences do not depend on the contracts' consequences in states of the world in which they are identical. In this case the no-trade result does not depend on any homogeneity in prior beliefs, and, in fact, the theorem is formalized and proved without any probabilistic reasoning. Furthermore, we show that

¹For discussions on this subject see Aumann (1987), Bernheim (19865) and Morris (1995)

the no-trade result in a noiseless information structure is a special case of the Agreement Theorem. The intuition behind this result is that if it is common knowledge that all agents want to trade at a given state, they could have done so ex-ante, so that the existing allocation could not be ex-ante efficient. As summarized by Geanakoplos (1992): “Common knowledge of rationality and of optimization eliminate trade”. This section may be viewed as an abstraction of Milgrom and Stokey (1982) result, for the simplified case of noiseless signals. This enables us to present the mathematical beauty of their No-Trade Theorem.

Next, we analyze the interaction between agents whose preferences do not conform to the Sure Thing Principle. We demonstrate that although superficial analysis may indicate that trade is possible, it is really a result of a basic flaw at the individual rational level, i.e. *dynamic inconsistency*. In Appendix B we give a simple example in which trade between such agents may be accomplished. The agents’ utility functions are not linear in probabilities, i.e., do not satisfy the *Independence Axiom*, which is the Sure Thing Principle’s parallel in a world of objective probabilities. Technically, the trading result is possible since the absence of the Independence Axiom, information can change the agents’ marginal rate of substitution between states of the world which they think are possible. In the example we give, information is symmetric, and trade is driven not by difference in signals, but from dynamic inconsistency of the single individual. Dynamic consistency is the weakest form of rationality which one could impose in a dynamic framework. Hence, any analysis which asserts rationality, has to suppose it at the outset.

Following the works of Hammond (1988), Machina (1989), Segal (1990), and Karni and Schmeidler (1991b), who assume objective probability, we investigate the relations between the *Sure Thing Principle* and *Dynamic Consistency* in a decision theoretic framework, but with no objective probabilities. We define a *consequentialist* decision maker as one whose conditional preference (given an event) does not depend on uncertainties that did not happen. We show that if a preference relation is dynamically consistent and consequentialist, then it satisfies the Sure Thing Principle. Therefore, if we wish to investigate the interaction among agents who do not abide by the Sure Thing Principle and are rational, the consequentialism assumption has to be relaxed. Nevertheless, the question upon which consequences do preferences depend, is left open.²

²Our analysis *does not* attempt to separate the beliefs regarding the (subjective) likelihood of events, from the

According to one approach, the agent should take into account what he would have done if he had reached the event which was not realized. The implication of this formulation is that at any event, the agent evaluates the whole contract, treating its impossible consequences and possible ones *symmetrically*. We refer to this type of preferences as *hypothetical non-consequentialist*. The conditional preference relation must be *complete*, and imposing dynamic consistency means it should conform to the ex-ante preferences. In other words, *there is no place for a natural evolution of the decision process as information is being revealed*. It comes with little surprise that dynamic consistency imposed on these preferences maintains the no-trade result.

As a result of the drawbacks of this formulation, and especially its deficiency in analyzing the effects of new information on the decision process, we advocated the decision theoretic approach, due to Machina (1989). According to this paradigm, the decision process is divided into two non-identical stages. Ex-ante, the decision maker chooses an optimal contract, which serves later as a “status quo” contract. Once the decision maker learns that an event has not happened, his conditional preferences at the event which did obtain depend on what would have happened in the non-realized event, *according to the original contract*. I.e., the contract which was a *real possibility* when those states were considered. Thus, when the decision maker learns that an event has occurred, he evaluates the contracts at the event in conjunction with the non-realized consequences of the status-quo contract. The consistency condition in the framework implies that the agent’s preferences over contracts at the event which happened, conform to his preferences ex-ante, taking borne uncertainties (at the event which did not obtain) as fixed.

We show that dynamic consistency alone might not be sufficient for the no-trade result for the status-quo non-consequentialist preferences. The main thrust behind this result is the lack of consequentialism: an agent cares about his welfare in states of the world that might have happened, but did not. This enlarges the set of possible contracts, since an identical written contract is interpreted differently by agents with different information. The lack of consequentialism enables agents not to reveal all their information while trading.

valuation of consequences associated with the realization of those events, in the agents’ preferences. Other works, as Epstein and Le Breton (1993) and Eichenberger and Kelsey (1996), which tried to impose such a separation, reached negative results regarding the compatibility of dynamic consistency and non-Bayesian models.

We formalize and prove a No-Trade Theorem for status-quo non-consequentialist preferences that are dynamically consistent, giving sufficient conditions for this case. In a last example we show that the non-neutrality of uncertainty (i.e., uncertainty aversion or love), combined with asymmetries in information, may induce trade between dynamically consistent, status-quo non-consequentialist agents, even when information is noiseless.

The discussion in the paper illuminates both pros and cons of non-expected utility paradigms. On the one hand, examining these general theories enabled us to shed new light on the no-trade result, and provide the minimal sufficient conditions for it. Under non-expected utility, we might have even uncovered a new source for trade as a result of asymmetric information. On the other hand, if we wish to maintain rationality (dynamic consistency) in a world where people are not expected utility maximizers, we are forced to give up very appealing axioms in decision theory and economics, such as consequentialism. Relaxing this assumption is not straightforward, and may yield ambiguous economic results, depending on alternative assumptions concerning the decision process.

2. The No-Trade Theorem and the Sure Thing Principle

In this section we prove the No-Trade Theorem for a noiseless information structure. Let $I = \{1, \dots, m\}$ be a finite set of agents. $\Omega = \{\omega_1, \dots, \omega_n\}$ is a finite state space such that every state is a complete description of the world (of the physical environment, knowledge, actions and preferences of all agents). Let Σ be the algebra of events on Ω . Every agent i is endowed with a knowledge function k^i on Ω such that $k^i(\omega)$ represents all the information agent i has at state ω .³ Let $\Pi^i(\omega) = \{\omega' \in \Omega : k^i(\omega') = k^i(\omega)\}$ be the set of all states which are indistinguishable for agent i from state ω . These classes of states (cells) constitute an information partition of the agent (collection of disjoint sets whose union is Ω). An event H obtains (happens) at ω if and only if $\omega \in H$. Agent i knows event H at ω if $\Pi^i(\omega) \subseteq H$, i.e., at all states i thinks possible at ω , H obtains. The event “ i knows H happened” is the set of states at which i knows H obtained: $\{\omega : \Pi^i(\omega) \subseteq H\}$, i.e., an event is self-evident if at the time of happening all agents know it. An event B will be *Common Knowledge*

³Alternatively, we could think of $k^i(\omega)$ as a signal the agent receives when the state is ω . The assumption that k^i is a function (and not a correspondence) is equivalent to assuming that the information (signal) is without noise.

among I at ω if all agents know B , all agents know that all agents know B , all agents know that all agents know that all know B , etc. Aumann (1976) proved that this complex infinite sequence of conditions is equivalent to the single condition that there exists a self-evident event E such that $\omega \in E$ and $E \subseteq B$.

Let C be the set of consequences.. A contingent contract is a function from Ω to C . Let A be the set of all contingent contracts. Define for every $E \in \Sigma$: $a =_E b$ if $a(\omega) = b(\omega)$ for every $\omega \in E$, and let $E^c = \Omega \setminus E$ be the complement of E . Assume (after Savage) for every $i \in I$:

Weak Order (P1): \succ^i over contingent contracts is weak order.⁴

Sure Thing Principle (P2): If $(a =_E a', b =_E b', a =_{E^c} b, a' =_{E^c} b')$ then $(a \succ^i b$ if and only if $a' \succ^i b')$.

The Sure Thing Principle asserts that preference between two contracts should not depend on those states in which they have the same consequences.⁵ Given P2 it is natural to define conditional preferences in the following way:

Definition 1: Let $a, b \in A$ and $E \in \Sigma$: $a \succ_E^i b$ (a is preferred to b given E) if and only if $a' \succ^i b'$ for all $a', b' \in A$ such that: $(a =_E a', b =_E b', a' =_{E^c} b')$.

Given P2 it is sufficient to find one pair (a', b') for which $a' \succ b'$. The following Lemma justifies the use of the above definition of conditional preferences in conjunction with the Sure Thing Principle.

Lemma 1: *If \succ^i satisfies P1 and P2 then \succ_E^i is a preference relation.*

Proof: Appendix A.

The following property of conditional preference relation relates the agent's preferences given disjoint events to his preferences given their union.

Definition 2: A preference relation \succ^i over A satisfies the *Conditional Dominance Principle* (CDP) if for every $E_1, E_2 \in \Sigma$ such that $E_1 \cap E_2 = \emptyset$ and $a, b \in A$, if $a \succ_{E_1}^i b$ and $a \succ_{E_1 \cup E_2}^i b$.

Lemma 2: *If \succ^i over A is a weak order then preferences abide by the CDP.*⁶

⁴Alternatively, \succ^i is an asymmetric and negative transitive binary relation, or a preference relation.

⁵An additive across-states representation of preferences: $U(a) = \sum_{\omega \in \Omega} u[a(\omega), \omega]$ abides by the Sure Thing Principle. Furthermore, if the given preferences are complete, transitive, continuous and the preferences abide by P2, they are additive across states.

⁶Although the Sure Thing Principle (P2) is not needed in the proof, it appears in the lemma because only if

Proof: Appendix A.

Thus, if preferences satisfy P1 and P2, the conditional preference relation given any event is not empty, and satisfies the Conditional Dominance Principle.

Definition 3: A contingent contract b is *ex-ante efficient* if and only if there does not exist another contingent contract a such that $a \succ^i b$ for all agents $i \in I$.⁷

Theorem 1: *Let all agents have preferences over the set of contingent contracts A which satisfy P1 and P2. If b is an ex-ante efficient contract, there does not exist a state $\omega^* \in \Omega$ and another contingent contract a such that the event: $D = \{\omega: a \succ_{\Pi^i(\omega)}^i b \text{ for all } i \in I\}$ is common knowledge at ω^* .⁸*

Proof: Assume such ω^* exists. Let $E \subseteq D$ be self-evident and $\omega^* \in E$. Since D is common knowledge at ω^* , then for all $i \in I$ and $\omega' \in E$: $a \succ_{\Pi^i(\omega')}^i b$. E is the disjoint union of cells (for each agent) on which a is preferred to b . By the Conditional Dominance Principle $a \succ_E^i b$ for all $i \in I$.

Define: $a' = \begin{cases} a & \omega \in E \\ b & \omega \in E^c \end{cases}$. By the definition of conditional preferences, $a' \succ^i b$ for all $i \in I$, which contradicts the assumption that b was ex-ante efficient.

Thus, if information is noiseless, and agents' preferences abide by the Sure Thing Principle, they can not agree on a trade which will improve the welfare of every agent given his private information. In Appendix C we prove that the No-Trade Theorem may be viewed as a special case of Aumann's Agreement Theorem. In the following section we shall see that when preferences do not satisfy Savage's P2 agents may agree to disagree on the interpretation of a contract, which will result in the possibility of improving trade.

preferences satisfy P2, will the conditional preference as defined not be an empty binary relation.

⁷This is a weak definition of ex-ante efficiency. The stronger version - there does not exist a contract a such that $a \succ^i b$ for all agents and there exists at least one agent for whom $a \succ^i b$ - would not change any of the results.

⁸Dow *et al.* (1990) prove that in an exchange economy (like Milgrom-Stokey's) if markets are complete the no-trade theorem does not depend on the common prior. However, their result is a corollary of Milgrom and Stokey's original result who assume concordant belief. Assuming every payoff-relevant state has a single signal, or information is noiseless, is a stronger assumption on the priors. Moreover, both (Dow *et al.*, 1990 and Milgrom & Stokey, 1982) do not specify what conditions the decision makers should abide by for the no-trade result to hold. Here, we do not need even a representation by a utility function, only a preference relation which satisfies the Sure Thing Principle.

3. (No) Trade Without The Sure Thing Principle⁹

In Appendix B we analyze the interaction between agents whose utility is not linear in probabilities, i.e., does not satisfy the *Independence Axiom*, which is the Sure Thing Principle equivalent in a world of objective probabilities. We give a simple example in which trade between such agents may be accomplished as a pure informational phenomenon. This is so because in the absence of the Independence Axiom, information might change agent's marginal rate of substitution between states of the world which they think are possible. Although we technically maintain common knowledge of rationality, there is a basic flaw at the individual rationality level. In the example we give, information is symmetric (and even the priors are equal), and trade is driven not by difference in signals, but from dynamic inconsistency of the single individual. Ex-ante he prefers one contract over the other, while after the arrival of new information his preferences are reversed. Putting this in an atemporal dynamic framework (time here is negligible), this trade result is an extreme example of a pattern of behavior of non-expected utility maximizers, which has been widely criticized as being dynamically inconsistent (Machina, 1989, gives a detailed survey of the subject).

3.1. Dynamic Consistency and Trade

An example could clarify the idea of dynamic consistency. It relies on an example given by Ellsberg (1961) to show the difference between attitudes toward risk and uncertainty. Let there be two urns, each containing 100 balls, which can be either red or black. It is known that one urn holds 50 red and 50 black balls. The number of red (black) balls in the other urn is unknown. Two balls are drawn at random, one from each urn. The subject is asked to bet on the color of one of the balls. A correct guess wins \$100; incorrect guess loses nothing (and pays nothing). Ellsberg found that a non-negligible part of the population has the following preferences:

Indifferent between betting on red or black from the first urn.

Indifferent between betting on red or black from the second urn.

Prefers to bet on red drawn from the first urn to red drawn from the second urn.

Prefers to bet on black drawn from the first urn to black drawn from the second urn.

⁹Readers who are not familiar with non-expected utility are referred to Karni and Schmeidler (1991a).

This pattern of preferences is inconsistent with the Sure Thing Principle since the decision maker prefers objective risk to uncertainty (the expected utility paradigm is not rich enough to differentiate between the two). Nevertheless, his preferences may be represented by an expected utility functional with respect to a non-additive probability measure (Gilboa, 1987 and Schmeidler, 1989) or a maximum expected utility with respect to a non-unique prior (Gilboa and Schmeidler, 1989).

Now, suppose a decision maker is told that the two balls drawn were red and black, but not which ball was drawn from which urn. He is asked to bet as before, which is equivalent to betting on which urn contained the red (black) ball drawn. Suppose he prefers to bet that the red ball was drawn from the second urn. This is an inconsistent choice: before knowing the colors of the two balls drawn, he preferred to bet “red from the first urn” rather than “red from the second urn”. The two bets differ only in the event which happened (drawing two different balls) and equal elsewhere. Therefore the bettor had no reason to reverse his preferences. If he prefers to bet that the red ball was drawn from the first urn, it is equivalent to betting that the ball drawn from the second urn was black. This, again, is inconsistent with his preferences ex-ante, since he preferred “black from the first urn” over “black from the second urn”. The two bets are equal, except for the event that materialized. Indifference would not help here, because consistency of preferences would imply the subject should have been indifferent between the bets ex-ante, too.

The conclusion drawn from this example is that if we wish to maintain dynamic consistency of preferences that do not satisfy the Sure Thing Principle, preferences over contracts should depend on events that did *not* obtain. Nevertheless, the question upon which consequences do preferences depend, is left open.¹⁰ According to one approach the agent should take into account what he would have done if he had reached the event which was not realized. The implication of this formulation is that at any event, the agent evaluates the whole contract, treating its possible and impossible consequences *symmetrically*. The conditional preference relations must be *complete*, and imposing dynamic consistency means it should conform to the ex-ante preferences. In other words, *there is no place for a natural evolution of the decision process as information is being revealed*. It comes with little surprise that dynamic consistency imposed on this preferences maintains the no-trade

¹⁰The presentation of the ambiguity was triggered by discussions with Luca Anderlini and George Mailath.

result.¹¹

As a result of the drawbacks of the above formulation, and especially its deficiency in analyzing the effects of new information on the decision process, we advocated the decision theoretic approach, due to Machina (1989). According to this paradigm, the decision process is divided into two non-identical stages. Ex-ante, the decision maker chooses a bet (contract), which serves later as a “status quo” bet. Once the decision maker learns that an event has not happened, his conditional preferences at the event which did obtain, depend on what would have happened in the non-realized event, *according to the original contract*. I.e., the contract which was a *real possibility* when those states were considered. Thus, when the decision maker learns that an event has occurred, he evaluates the bets (contracts) at that event in conjunction with the non-realized consequences of the status quo bet. Following, we formally define these notions.

For every $E \in \Sigma$ let $a|_E$ be the *contract a at E* such that for every $\omega \in E$: $a|_E(\omega) = a(\omega)$. Of course, two different contracts could be equal at E. For every $a, b \in A$ let the *composite contract a at E given b* be: $b/a|_E = \begin{cases} a & \omega \in E \\ b & \omega \in E^c \end{cases}$. Let \succ_E^i be the preference relation of agent i on contracts at E.¹²

DEFINITION 4: Let b be a *status quo contract* for agent i , i.e., a contract the agent chose or was endowed with when no information was available (at -). The agent’s preferences over contracts at E will be *status-quo non-consequentialist* if for any two contract $a, a' \in A$: $a \succ_E^i a'$ if and only if $b/a|_E \succ_E^i b/a'|_E$.

We maintain the same notion of ex-ante efficiency as in Definition 3.

DEFINITION 5: Preference relations $\{\succ_E^i\}_{E \in \Sigma}$ over contingent contracts satisfy *Dynamic Consistency* if for all $E, F \in \Sigma$ such that $E \subseteq F$ and for all contracts $a, b \in A$ such that $(a =_{E^c} b)$: $(a \succ_E^i b)$ if and only if $(a \succ_F^i b)$.

In Figure 1 decision nodes appear in boxes, and contracts at events appear in circles. Heavy lines correspond to choices the decision maker would make, if he is offered a choice between sub-contracts. Thus, if his preferences are dynamically consistent, he prefers $a|_F$ over $b|_F$ (on the right

¹¹A formal treatment of these hypothetical non-consequentialist preferences is deferred until section 4.

¹²If $E = -$ we omit the “at - ” subscript.

hand side) if and only if he would choose $a|_E$ over $b|_E$ (in the left tree).

After defining formally dynamic consistency, we are left with the question whether imposing this condition is sufficient to preclude trade as an informational phenomenon. The following example shows why it is not:

Assume two states of the world, $\Omega = \{\omega_1, \omega_2\}$, and two agents: Alice (A) and Bob (B). The information structures are: $\Pi^A = \{(\omega_1), (\omega_2)\}$ and $\Pi^B = \{(\omega_1, \omega_2)\}$, i.e., after a state of nature has occurred, Alice knows it while Bob remains ignorant. Assume contract b is ex-ante efficient. It may well be that once Alice knows the true state, she would prefer switching to contract a (since her preferences may depend on unrealized states and she is status-quo non-consequentialist), and therefore by dynamic consistency those were her preferences ex-ante, too: $(b, a) \succ^A (b, b)$ and $(a, b) \succ^A (b, b)$ but $(b, b) \succ^A (a, a)$, i.e., ex-ante she prefers contract b over a . For Bob: $(b, b) \succ^B (b, a)$ and $(b, b) \succ^B (a, b)$ (otherwise b would not have been ex-ante efficient) but $(a, a) \succ^B (b, b)$ i.e., ex-ante (and for Bob, ex-post as well) Bob prefers a to b . Whatever state of the world is realized, both agents will find it profitable to opt for contract (a, a) . Since Alice has borne part of the uncertainty with contract b , for her the contract would be (a, b) or (b, a) and if Bob has the above preferences he would agree. Formally, the event: $D = \{\omega: b/a|_{\Pi(\omega)} \succ_{\Pi^i(\omega)}^i b \text{ for all } i \in I\} = \{\omega_1, \omega_2\} = \Omega$ is common knowledge at any state.¹³

3.2. Dynamic Consistency and The Sure Thing Principle

The object of this subsection and the following, is to define sufficient conditions for the No-Trade Theorem for preferences that do not abide by the Sure Thing Principle, and are status-quo non-consequentialist. To this aim, we first define formally consequentialism, and investigate the relation between dynamic consistency and the Sure Thing Principle in the presence of consequentialism.

DEFINITION 6: Preference relations $\{\succ_E^i\}_{E \in \Sigma}$ over contingent contracts satisfy *consequentialism* if for all $E \in \Sigma$ and all $a, a', b, b' \in A$ such that $(a =_E a', b =_E b', a =_{E^c} b, a' =_{E^c} b')$: $a \succ_E^i b$ if and only if $a' \succ_E^i b$.

¹³Note that Alice's preferences over written contracts have evolved as a result of the arrival of new information. On the Dutch book potential in the status-quo non-consequentialist preferences, which is present in the above example, see Machina's (1989) discussion on Segal's and Dekel's examples.

Consequentialism implies that preferences over sub-contracts are independent of previous uncertainties. Given an event, a decision maker acts as if he had started out from that point and treats the uncertainty involved in the preceding part of the tree as irrelevant or as if it never existed. It is apparent that consequentialism is closely related to the Sure Thing Principle, although the former is a characteristic of preference at any event while the latter relates to preference at - .¹⁴ The following proposition relates these three concepts.

PROPOSITION 1: *If preferences over contracts $\{\succ_E^i\}_{E \in \Sigma}$ satisfy both Dynamic Consistency and Consequentialism then preferences over contingent contracts satisfy the Sure Thing Principle.*¹⁵

PROOF: Take $a, b \in A$ such that $a =_E a', b =_E b', a =_{E^c} b, a' =_{E^c} b'$: $a \succ^i b$ if and only if (by dynamic consistency) $a \succ_E^i b$ if and only if (by consequentialism) $a' \succ_E^i b'$ if and only if (by dynamic consistency again) $a' \succ^i b'$.

As noted earlier, unlike Epstein and Le Breton (1993) we do not try to impose any separation between beliefs and valuation of consequences. Any attempt to impose Machina and Schmeidler's (1992) "probabilistic sophistication" (by strengthening Savage's P4) will result in convergence to the Bayesian paradigm. The latter is not rich enough to represent the decision maker's vague belief, in the presence of uncertainty.

3.3. No Trade without The Sure Thing Principle

An immediate conclusion of Proposition 1 is that dynamic consistency and consequentialism are sufficient for the no-trade result. However, imposing both does not allow us to represent preferences which do not abide by the Sure Thing Principle. Next, we introduce a weaker condition which, together with dynamic consistency, is sufficient for the no-trade theorem, for preferences which are

¹⁴Machina (1989) claims that when preferences do not satisfy the independence axiom one should give up either dynamic consistency or consequentialism in order to maintain the other. He argues in favor of abandoning consequentialism because of the non-separability of preferences when independence is not maintained.

¹⁵Note that consequentialism and the Sure Thing Principle are not sufficient for dynamic consistency as the latter is the axiom which connects ex-ante to interim decisions, while the former are concerned with ex-ante separately. This is similar to the case of objective probabilities, where a third axiom was needed (reduction of compound lotteries). However, once we take the definition of conditional preferences as a definition of interim preferences, dynamic consistency and consequentialism become necessary for the Sure Thing Principle too.

status-quo non-consequentialist. Trade resulting from informational incentives will be impossible to accomplish, for agents whose preferences satisfy this condition, although they do not conform to the Sure Thing Principle.

DEFINITION7: Preference relations $\{\succ_E^i\}_{E \in \Sigma}$ over contingent contracts satisfy *The Generalized Conditional Dominance Principle (GCDP)* if for all $a, b \in A$, $E \in \Sigma$ and a partition E_1, \dots, E_n of E : If $b \setminus a|_{E_j} \succ_{E_j}^i b$ for every $j = 1, \dots, n$ then $b \setminus a|_E \succ_E^i b$.

This principle is a generalization of the conditional dominance principle for status-quo non-consequentialist preferences. It says that given a certain contingent contract b , if it is beneficial to deviate from b to a at each member of a set of disjoint events then it is advantageous to deviate from b to a at their union. Theorem 2 proves that the above condition, together with dynamic consistency, are sufficient to imply the no-trade result.

THEOREM 2: *Let all agents have status-quo non-consequentialist preferences $\{\succ_E^i\}_{E \in \Sigma}$ over contingent contracts at events, which satisfy Dynamic Consistency and the Generalized Conditional Dominance Principle. If b is an ex-ante efficient contract, then there does not exist a state ω^* and another contract a such that the event: $D = \{\omega: b \setminus a|_{\Pi^i(\omega)} \succ_{\Pi^i(\omega)}^i b \text{ for all } i \in I\}$ is Common Knowledge at ω^* .*

PROOF: Assume such w^* exists. Let $E \subseteq D$ be self-evident and $w^* \in E$. Since D is common knowledge at w^* then for all $i \in I$ and $\omega' \in E$: $b \setminus a|_{\Pi^i(\omega')} \succ_{\Pi^i(\omega')}^i b$. Since E is self-evident, it is the disjoint union of events (for each agent) at which the agent prefers to deviate to a . By *GCDP*: $b \setminus a|_E \succ_E^i b$. By dynamic consistency: $b \setminus a|_E \succ_E^i b$ for all $i \in I$, which contradicts the assumption that b was ex-ante efficient. Q.E.D.

3.4. Trade without The Sure Thing Principle

In this section we give a numerical example of a situation in which the agents' preferences are not uncertainty neutral (and hence do not abide by the Sure Thing Principle). We assume status-quo non-consequentialist preferences and impose dynamic consistency. We show that asymmetric information may lead to trade which is commonly known to dominant the ex-ante efficient allocation for every agent. Thus, dynamic consistency of preferences is not sufficient for a no-trade result for the class of non-additive expected utility preferences, if they are status-quo non-consequentialist.

The following example builds upon the example given in [Elsberg \(1961\)](#), which was described above. The two urns are labeled I and II respectively. There are four possible states of the world: $\omega = RR, BR, RB, BB$ (read: Red from urn I and Red from urn II etc.). Suppose aggregate uncertainty in the world depends solely on the second urn. If the ball drawn from the second urn is red, there are \$300 to split between the agents. Otherwise there are only \$150 to allocate. The two agents, A (Alice) and B (Bob), are both risk-neutral but not uncertainty-neutral. Their non-additive probability measures are given in [Table I](#).

It can easily be shown that since Alice is more uncertainty averse about the relevant event on which the aggregate uncertainty depends, she is insured in every ex-ante efficient allocation (i.e. receives the same payoff in all states of the world), while Bob bears the burden of the uncertainty.¹⁶ Take the ex-ante efficient allocation: $b^A = (100, 100, 100, 100)$ and $b^B = (200, 200, 50, 50)$. Let $\Pi^A = \{(RR, RB), (BR, BB)\}$, i.e., Alice observes the color of the ball drawn from the first urn only, which resolves the risk (urn I), but not the uncertainty. Bob receives no information (i.e., Π^B is the trivial partition of ω).

We claim that in every state of nature b is not interim Pareto-efficient, i.e, there exists a trade option that Pareto dominates b for both Alice and Bob, given their private information. E.g., let $a^A = (200, 200, 50, 50)$ and $a^B = (100, 100, 100, 100)$, so they will actually switch roles relative to the original contract. It can easily be calculated that: $U^A(b^A) = 100$

$$U^A(b^A/a^A_{\{\{RR, RB\}\}}) = U^A(200, 100, 50, 100) = U^A(b^A/a^A_{\{\{BR, BB\}\}}) = U^A(100, 200, 100, 50) = 102.5$$

$$U^A(a^A) = U^A(200, 200, 50, 50) = 87.5$$

$$U^B(b^B) = 95$$

$$U^B(100, 200, 100, 50) = U^B(200, 100, 50, 100) = 94.5$$

$$U^B(a^B) = U^B(100, 100, 100, 100) = 100.$$

Before knowing the result of urn I, Alice prefers b^A to a^A because a^A offers only \$50 if the ball drawn from the second urn is black, and this event has relatively “high” weight (since she is uncertainty averse in that region). But once she learns the result of urn I, her preferences over contracts at the

¹⁶The proof that Alice is insured in every ex-ante efficient allocation could be given in two ways: directly (this is a standard linear programming problem), or by using [Proposition B.1](#) and the relation between [Yaari’s dual theory](#) and non-additive expected utility theory ([Karni and Schmeidler, 1991a](#)).

relevant events are as if she already consumed \$100 in probability 0.5, and she is now willing to take an uncertain stand so as to win more \$100 or to lose \$50. This is a result of the non-linearity of the measure σ^A . Bob will opt to the suggested contract since his uncertainty aversion is such that he prefers to be insured over the original contract. Therefore, if preferences are represented by an expected utility functional with respect to a non-additive probability measure, agents could be dynamically consistent, but information may induce trade between them.

4. HYPOTHETICAL NON-CONSEQUENTIALISM AND FULL DYNAMIC CONSISTENCY

In this section we analyze hypothetical non-consequentialist preferences (i.e. the conditional preferences do not refer to the status-quo contract) which were introduced in section 3. As noted there, the agent takes into account what he would have done if he had reached the event which was not realized. We show that dynamic consistency of these preferences are sufficient for the no-trade result, as it abolishes any opportunity for the evolution of preferences as information is revealed to the agent.

DEFINITION 8: Preference relations $\{\succ_E^i\}_{E \in \Sigma}$ over contingent contracts satisfy *Full Dynamic Consistency* if for all $E, F \in \Sigma$ such that $E \subseteq F$ and for all contracts $a, b \in A$: $(a \succ_E^i b)$ if and only if $(a \succ_F^i b)$.

The full dynamic consistency is stronger than definition 5 since it does not restrict the contracts to be equal on the complement of E. However, this is the appropriate definition for preferences which are hypothetical non-consequentialist.

PROPOSITION 2: *Let all agents have preferences $\{\succ_E^i\}_{E \in \Sigma}$ over contingent contracts at events which are hypothetical non-consequentialist and satisfy Full Dynamic Consistency. If b is an ex-ante efficient contract, then there does not exist a state ω^* and another contract a such that the event: $D = \{\omega: a \succ_{\Pi^i(\omega)}^i b \text{ for all } i \in I\}$ is Common Knowledge at ω^* .*

PROOF: Assume that such state and contract exist. Then, by strong dynamic consistency all agents prefer a to b ex-ante, in contradiction to the assumption that b was ex-ante efficient. Q.E.D.

5. CONCLUSION

This paper focuses attention on the minimal sufficient restrictions on the decision makers preferences for the existence of a No-Trade Theorem, in an economy where information is noiseless and hence markets are complete: a very basic notion of dynamic consistency, the generalized conditional dominance principle and the strong assumption of common knowledge of rationality. Thus, the main conclusions of this paper is that if one believes that behavior satisfies the Sure Thing Principle, then the explanation of trade as an informational phenomenon must concentrate on the lack of common knowledge of rationality, at some level of the hierarchy of knowledge (see, e.g., Neeman, 1996), or on the missing markets and lack of a common prior (Morris, 1994). When the Sure Thing Principle is relaxed, we showed that for hypothetical non-consequentialist preferences, dynamic consistency was a sufficient condition for information not to induce trade. In the more interesting case of status-quo non-consequentialist preferences, trade could be explained by the non-neutrality of uncertainty, even if information is noiseless, let alone when signals are noisy.

Dept. of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104-6297, U.S.A.; halevy@econ.sas.upenn.edu

6. APPENDIX A

6.1. Proofs

LEMMA 1: *If \succ satisfies P1 and P2 then \succ_E is a preference relation (weak order).*

PROOF: Asymmetry, i.e., if $a \succ_E b$ then $\sim (b \succ_E a)$. Assume $b' \succ_E a$. Then there exist $a', b' \in A$ such that $(a =_E a', b =_E b', a' =_{E^c} b')$ and $b' a'$. Contrary to the assumption that for every two contracts of the above structure $a' b'$ (asymmetry of \succ). Negative transitivity, i.e., if $\sim (a \succ_E b)$ and $\sim (b \succ_E c)$ then $\sim (a \succ_E c)$. If $\sim (a' \succ b')$ then there exist $a', b' \in A$ such that $(a =_E a', b =_E b', a' =_{E^c} b')$ and $\sim (a' b')$. If $\sim (b \succ_E c)$ then there exist $b'', c'' \in A$ such that $(b =_E b'', c =_E c'', b'' =_{E^c} c'')$ and $\sim (b'' \succ c'')$. Define: $(a'' =_E a, a'' =_{E^c} c'')$. Note that $(a' =_E a'', b'' =_E b', a' =_{E^c} b', a'' =_{E^c} b'')$ and by assumption $\sim (a' b')$. Hence by P2: $\sim (a'' \succ b'')$. By assumption $\sim (b'' \succ c'')$; hence, because is negative transitive: $\sim (a'' \succ c'')$, i.e. $\sim (a'' \succ_E c)$. Q.E.D.

LEMMA 2: *If \succ over A is a weak order then preferences obey by the CDP.*

PROOF: Let $E_1, E_2 \in \Sigma$ such that $E_1 \cap E_2 = \emptyset$ and $a, b \in A$ such that $a \succ_{E_1} b$ and $a \succ_{E_2} b$.

Let $c \in A$ be any other contract. Define:

$$a' = \begin{cases} a & \omega \in E_1 \cup E_2 \\ c & \omega \in (E_1 \cup E_2)^c \end{cases} \quad d = \begin{cases} b & \omega \in E_1 \\ a & \omega \in E_2 \\ c & \omega \in (E_1 \cup E_2)^c \end{cases} \quad b' = \begin{cases} b & \omega \in E_1 \cup E_2 \\ c & \omega \in (E_1 \cup E_2)^c \end{cases}$$

By the definition of conditional preference $a' \succ d$ and $d \succ b'$ then by P1 $a' b'$. Since c was arbitrary, by the definition of conditional preference: $a \succ_{E_1 \cup E_2} b$. Q.E.D.

7. APPENDIX B

THE PATTERN OF TRADE UNDER THE DUAL THEORY OF CHOICE UNDER RISK

The expected utility hypothesis assumes that agents' utility from a contingent contract is given by their expected utility (with respect to an additive probability measure) from the contract in different states. Because expectation is a linear operator, the additional information does not change the marginal rate of substitution between states which agents consider possible. In this Appendix we assume, for simplicity, that there is an objective probability distribution over states of the world (e.g., by lottery) which is known to all agents, and that trade is contingent on this stochastic process. Agents do not know the exact lottery result. Each agent is told that the outcome is a member of a set of possible outcomes (i.e., $\Pi^i(\omega)$). Assume that agents' preferences over possible contingent contracts are represented by Yaari's (1987) dual utility functional and that agents are risk-averse. This simplification is made to capture the general feature that utility is not linear in probabilities (the constant marginal utility of wealth plays no role here). Following, we shall present an example, with three states of the world, two agents with symmetric information but non-expected utility functions, who find it profitable to trade out of an ex-ante efficient allocation, after the arrival of new information.

7.1. The Ex-Ante Efficient Allocation

We focus first on the case where there are only two states of the world, $\omega = \{\omega_1, \omega_2\}$. The (objective) probability of state ω_1 is p . The utility of agent i from contract b , which promises him b_1^i and b_2^i in states ω_1 and ω_2 respectively, is given by:

$$U^i(b_1^i, p; b_2^i, 1-p) = \begin{cases} b_1^i[1 - f^i(1-p)] + b_2^i f^i(1-p) & b_1^i < b_2^i \\ b_1^i f^i(p) + b_2^i[1 - f^i(p)] & b_1^i > b_2^i \end{cases}$$

The marginal rate of substitution between b_2^i and b_1^i is given by:

$$MRS_{b_2, b_1}^i = \begin{cases} \frac{1-f(1-p)}{f(1-p)} & b_1^i < b_2^i \\ \frac{f(p)}{1-f(p)} & b_1^i > b_2^i \end{cases}$$

Note that $f(p)/(1-f(p)) < p/1(1-p) < (1-f(1-p))/f(1-p)$ as long as f is not the identity function. I.e., the marginal rate of substitution is higher/lower than the odds ratio. This result holds on the general class of EURDP as well, in the neighborhood of $b_1^i = b_2^i$.

DEFINITION B.1: Agent 2 will be called *Locally More Risk Averse* than agent 1 if the latter's willingness to pay for any simple lottery, of the structure $(x, 1-p; y, p)$ such that $x < y$, is higher than of the former's.

The above definition is a natural weakening of Yarri's (1986) definition of globally more risk averse under the Dual Theory. In the Dual Theory risk aversion is represented by the transformation function $f(\bullet)$ being increasing and convex. Agent 2 is *globally more risk averse* than agent 1 if there exist a convex transformation g such that $f^2(p) = g[f^1(p)]$ for every p . If agent 2 is more risk averse than agent 1, then $f^1(p) \geq f^2(p)$ for every p . Since in the Dual Theory the utility of a random variable equals its certainty equivalent, this condition suggests that the amount of money agent 1 is willing to pay in order to purchase any simple lottery ticket is at least as large as the amount of money agent 2 is willing to pay for the same ticket.

Next, we analyze the efficient allocations when there are only two agents. Assume $\bar{b}_1 > \bar{b}_2$ (aggregate uncertainty). Agent i is called insured if $b_1^i = b_2^i$.

PROPOSITION B.1: *Agent 2 is insured in an efficient allocation if and only if he is locally more risk averse than agent 1 at p .*

PROOF: Agent 2 is insured in an efficient allocation if and only if $MRS_{b_2}^1, b_1(b_1^1 > b_2^1) > MRS_{b_2}^2, b_1(b_1^2 > b_2^2)$ if and only if $f^1(p)/(1-f^1(p)) > f^2(p)/(1-f^2(p))$ which holds if and only if $f^1(p) > f^2(p)$, i.e. agent 2 is more risk averse than agent 1 at p . Q.E.D.

7.2. Trade as an Informational Phenomenon

The model we present here is minimal but sufficient to show how the arrival of new information may be an incentive to trade out of the ex-ante efficient contract. This example relies on the dynamic choice of each single agent (defined formally in section 3) and not on the asymmetric information.¹⁷

¹⁷A similar result appears in Dow *et al.* (1990). Their result, as noted by Epstein and Le Breton (1993) and as our present example illustrates, relies merely on dynamic inconsistency of the individual agents. Their claim that trade is a result of asymmetric information is not accurate: we show below that it could be reached with completely

Let there be three states (not in Savage's terminology, but in the one used above, i.e., the contracts are contingent on a roulette lottery in which outcomes are known only at the time of the trade): $\omega = \{\omega_1, \omega_2, \omega_3\}$. There exists an objective probability for the three states which is common knowledge among the agents $p = (p_1, p_2, p_3)$. Assume two agents $I = \{1, 2\}$ with symmetric information structures: $\Pi_1 = \Pi_2 = \{(\omega_1, \omega_2), (\omega_3)\}$. Let the contingent contracts have the same structure as before, i.e., let b_j^i be the payment to the i -th agent in the j -th state, $b_j^i \geq 0$, and $\sum_i b_j^i = \bar{b}_j$. Both agents have preferences over contingent contracts which are represented by a utility function as in the Dual Theory. Let j, k, l be a permutation of $\{1, 2, 3\}$ such that $b_j^i \leq b_k^i \leq b_l^i$. Then: $U^i(b) = b_j^i[1 - f^i(p_k + p_l)] + b_k^i[f^i(p_k + p_l) - f^i(p_l)] + b_l^i f^i(p_l)$. Assume the following probability transformations: $f^1(p) = (e^p - 1)/(e - p)$ and $f^2(p) = p^{1.8}$. For low values of p agent 2 is locally more risk-averse than agent 1 and for high values of p agent 1 is locally more risk-averse than agent 2. Let $1 = \bar{b}_3 = \bar{b}_2 < \bar{b}_1 = 2$ and $(p_1, p_2, p_3) = (0.4, 0.2, 0.4)$. In what follows we prove that in all the ex-ante efficient contracts agent 1 "insures" agent 2.

PROPOSITION B.2: *A contingent contract b^* will be ex-ante efficient if and only if $b_1^2 = b_2^2 = b_3^2$ and $b_1^1 > b_2^1 = b_3^1$.*

PROOF: States 2 and 3 are symmetric from the payoff perspective, so the agents will insure themselves in those states. The aggregate uncertainty in payoffs is between states $\{\omega_1\}$ and $\{\omega_2, \omega_3\}$. The problem reduces to the one studied in Proposition B.1, when the probability of the higher aggregate payoff is $p_1 = 0.4$. Therefore agent 1 insures agent 2. Q.E.D.

The event $D = \{\omega_1, \omega_2\}$ is common knowledge at w_1 , and the posterior probabilities at ω_1 and ω_2 are: $(\tilde{p}_1, \tilde{p}_2) = 2/3, 1/3$. Now, $f^1(2/3) < f^2(2/3)$. Therefore, agent 1 is locally more risk-averse than agent 2. At ω_1 and ω_2 , agent 2 could insure agent 1, leaving room for mutually beneficial trade, and both agents know and agree on the set of states in which there is another contract which dominates the no-trade. The event $D = \{w_1, w_2\}$ is common knowledge at w_1 , and the posterior probabilities at w_1 and w_2 are: $(p_1, p_2) = 2/3, 1/3$. Now, $f_1(2/3) < f_2(2/3)$. Therefore, agent 1 is locally more risk-averse than agent 2. At w_1 and w_2 , agent 2 could insure agent 1, leaving room for mutually beneficial trade, and both agents know and agree on the set of states in which there is

symmetric information and even with a common prior.

another contract which dominates the no-trade.

8. APPENDIX C

THE NO-TRADE THEOREM AS A SPECIAL CASE OF THE AGREEMENT THEOREM

In what follows we shall see that Theorem 1 is a special case of Aumann's Agreement Theorem. Let Z^i be agent i 's set of possible actions. A *Decision procedure* of agent i is a function from Σ to Z^i such that $D^i(H)$ is the recommendation of the decision procedure to agent i with information H .

DEFINITION C1: A decision procedure D^i satisfies the *Conditional Dominance Principle* (CDP) (or the *Sure Thing Principle for decision procedures*, or *Union Consistency*) if for every $E_1, E_2 \in \Sigma$ such that $E_1 \cap E_2 = \emptyset$ and $D^i(E_j) = z$ $j = 1, 2$ then $D^i(E_1 \cup E_2) = z$.

Assume that the decision procedures of all agents satisfy the CDP. An *Action Function* d^i is the action of agent i at state ω , or the function which implements D^i at ω . Formally: $d^i: \Sigma \rightarrow Z^i$ such that $d^i(\omega) = D^i(\Pi^i(\omega))$. Define the event " i takes action z^i " as: $\{d^i = z^i\} \equiv \{\omega \in \Sigma : d^i(\omega) = z^i\}$.

THE AGREEMENT THEOREM: *Let $z^i \in Z^i$ for all $i \in I$. If $\bigcap_{i \in I} \{d^i = z^i\}$ is common knowledge at ω , then there exists an event E such that $D^i(E) = z^i$ for every agent i .*

In words: If it is common knowledge at ω that agent i takes action z^i , then there exists an event E such that z^i is the recommendation of agent i 's decision procedure given E . Note that if $D^i \equiv D$ (the decision procedures are identical) then $z^i = z$ for all agents (all agents take the same action). The proof may be found in Aumann (1995).

To prove Theorem 1 as a corollary of the above theorem define: $Z^i = \{T, N\}$ (T for "Trade", N for "No trade"). Let b be an ex-ante efficient contract and let a be any other contract. The decision procedure of agent i is:

$$D^i(E) = \begin{cases} T & a \succ_E^i b \\ N & b \succ_E^i a \end{cases}.$$

If \succ^i satisfies Savage's P2 then D^i satisfies the CDP. Agents will trade b for a if and only if there exists a state ω^* at which it is common knowledge that all agents are willing to trade. By the Agreement Theorem, there exists an event E , such that $a \succ_E^i b$ for all agents. Define:

$a' = \begin{cases} a & \omega \in E \\ b & \omega \in E^c \end{cases}$. By the definition of conditional preference: $a' >^i b$ for all agents, which contradicts the assumption that b was ex-ante efficient.¹⁸¹⁹

Savage's P2 was sufficient for the decision procedure to satisfy the CDP. In section 3 we show it is not necessary, and investigate pattern of trade when preferences do not conform to this necessary conditions.²⁰

¹⁸The corollary is more general and claims that if b ex-ante dominates a for all agents there exists no state at which it is common knowledge that a dominates b for all agents given their information.

¹⁹Bacharach (1985) tried to prove the no-trade theorem using the agreement theorem. However, he assumes that all decision procedures are identical, i.e., that agents have identical preferences and share a common prior. Rubinstein and Wolinsky (199), using a similar methodology, assume a common prior.

²⁰Aumann (1995) notes that they should be related, but claims they are almost equivalent.

References

- [1] Anscombe, F. J., and R. J. Aumann: A Definition of Subjective Probability. *The Annals of Mathematics and Statistics* 34, 199-205 (1963).
- [2] Aumann, R. J.: Agreeing to Disagree. *The Annals of Statistics* 4, 1236-1239 (1976).
- [3] ———: Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica* 55 (1), 1-18 (1987).
- [4] ———: Interactive Epistemology. The Hebrew University of Jerusalem, Center for Rationality and Interactive Decision Theory, discussion paper #67 (1995).
- [5] Bacharach, M.: Some Extensions of a Claim of Aumann in an Axiomatic Model of Knowledge. *Journal of Economic Theory* 37, 167-190 (1985).
- [6] Bernheim, D.: Axiomatic Characterization of Rational Choice in Strategic Environment. *Scandinavian Journal of Economics* 88, 473-488 (1986).
- [7] Dow, J., V. Madrigal and S. R. C. Werlang: Preferences, Common Knowledge, and Speculative Trade. Fundacao Getulio Vargas Working Paper (1990).
- [8] Eichberger, J. and D. Kelsey: Uncertainty Aversion and Dynamic Consistency. *International Economic Review* 37 (3), 625-640 (1996).
- [9] Elsberg, D.: Risk, Ambiguity and The Savage Axioms. *Quarterly Journal of Economics* 75, 643-669 (1961).
- [10] Epstein, L. and M. Le Breton: Dynamically Consistent Belief Must Be Bayesian. *Journal of Economic Theory* 61, 1-22 (1993).
- [11] Geanakoplos, J.: Common Knowledge. *Journal of Economic Perspectives* 6 (4), 53-82 (1992).
- [12] Gilboa, I.: Expected Utility with Purely Subjective Non Additive Probabilities. *Journal of Mathematical Economics* 16, 65-88 (1987).

- [13] Gilboa, I. and D. Schmeidler: Maxmin Expected Utility with A Non-Unique Prior. *Journal of Mathematical Economics* 18, 141-153 (1989).
- [14] Hammond, P. J.: Consequentialist Foundation for Expected Utility. *Theory and Decision* 25, 25-78 (1988).
- [15] Harsanyi, J.: Games of Incomplete Information Played by ‘Bayesian’ Players I-III. *Management Science* 14, 159-182, 320-334, 486-502 (1967/8).
- [16] Kandel, E. and N.D. Pearson: Differential Interpretation of Public Signals and Trade in Speculative Markets. *Journal of Political Economy* 103 (4), 831-872 (1995).
- [17] Karni, E. and D. Schmeidler: Utility Theory with Uncertainty. *Handbook of Mathematical Economics* Vol. IV, Edited by W. Hildenbrand and H. Sonnenschein, Netherlands: Elsevier Science Publishes B.V. (1991a).1
- [18] ____: Atemporal Dynamic Consistency and Expected Utility Theory. *Journal of Economic Theory* 54, 401-408 (1991b).
- [19] Lang, L.H.P., R.H. Litzenger and V. Madrigal: Testing Financial Market Equilibrium under Asymmetric Information. *Journal of Political Economics* 100(2), 317-348 (1992).
- [20] Machina, M. J.: Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty. *Journal of Economic Literature* 27, 1622-1668 (1989).
- [21] Machina, M.J. and D. Schmeidler: A More Robust Definition of Subjective Probability. *Econometrica* 60 (4), 745-780 (1992).
- [22] Milgrom, P. and N. Stokey: Information, Trade and Common Knowledge. *Journal of Economic Theory* 26, 17-27 (1982).
- [23] Morris, S.: Trade with Heterogeneous Prior Beliefs and Asymmetric Information. *Econometrica* 62 (6), 1327-1347 (1994).
- [24] ____: The Common Prior Assumption in Economic Theory. *Economics and Philosophy* 11 (2), 227-253 (1995).

- [25] Neeman, Z.: Common Beliefs and the Existence of Speculative Trade. *Games and Economic Behavior* 16 (1), 77-96 (1996).
- [26] Rubinstein, A. and A. Wolinsky: On The Logic of 'Agreeing to Disagree' Type Results. *Journal of Economic Theory* 51, 184-193 (1990).
- [27] Savage, L. J.: The Foundations of Statistics. Second revised edition, New York: Dover Publications (1972).
- [28] Segal, U.: Two Stage Lotteries without the Reduction Axiom. *Econometrica* 58 (2), 349-377 (1990).
- [29] Schmeidler, D.: Subjective Probability and Expected Utility without Additivity. *Econometrica* 57, 571-587 (1989).
- [30] Yaari, M.E.: Univariate and Multivariate Comparisons of Risk Aversion: A New Approach. in: Uncertainty, information, and communication: Essays in honor of Kenneth J. Arrow, Volume III, edited by W.P. Heller, R.M. Starr and D.A. Starrett. Cambridge: Cambridge University Press (1986).
- [31] ____: The Dual Theory of Choice Under Risk. *Econometrica* 55, 95-115 (1987).