



Available online at [www.sciencedirect.com](http://www.sciencedirect.com)

SCIENCE @ DIRECT®

Games and Economic Behavior 46 (2004) 189–198

GAMES and  
Economic  
Behavior

[www.elsevier.com/locate/geb](http://www.elsevier.com/locate/geb)

# The possibility of speculative trade between dynamically consistent agents

Yoram Halevy

*The University of British Columbia, Department of Economics, #997-1873 East Mall,  
Vancouver, BC V6T 1Z1, Canada*

Received 8 August 2000

---

## Abstract

It is shown that interim dynamically consistent trade may be supported among agents who have resolute (non-consequential) choice preferences.

© 2003 Elsevier Inc. All rights reserved.

*JEL classification:* D81; D82; D84

---

## 1. Introduction

“No trade” theorems (Milgrom and Stokey, 1982; Holmström and Myerson, 1983) claim that if an initial contract is ex-ante Pareto efficient, asymmetric information cannot lead to trade among rational agents. It is often interpreted as referring to the impossibility of speculative trade (Geanakoplos, 1992). The no-trade results rely on two properties of Expected Utility preferences in dynamic choice problems: dynamic consistency and consequentialism. The existing literature (Dow et al., 1990) has demonstrated that failure of dynamic consistency may lead to trade, even when agents are symmetrically informed (Halevy, 1998). This paper shows that if agents are differentially informed, dynamically consistent agents may find it advantageous to trade at the interim stage.

As pointed out by Hammond (1988), in order to maintain dynamic consistency in a non-expected utility framework, it is necessary to relax other normative assumptions. This paper follows Machina (1989) in relaxing *consequentialism* (independence of

---

*E-mail address:* [yhalevy@interchange.ubc.ca](mailto:yhalevy@interchange.ubc.ca).

*URL:* <http://www.econ.ubc.ca/halevy>.

0899-8256/\$ – see front matter © 2003 Elsevier Inc. All rights reserved.  
doi:10.1016/S0899-8256(03)00044-7

counterfactual outcomes) in favor of a *resolute choice* (conditional preferences may depend on counterfactual events but are dynamically consistent) decision process.

The main contribution of the paper is to show, by means of an example, that dynamic consistency alone might not be sufficient for the no-trade result to hold, even in a noiseless information setting. It analyzes an environment in which trade occurs between two agents with resolute choice preferences. It shows that there may exist a commonly known event at which differentially informed agents, starting from an ex-ante Pareto efficient contract, are willing to exchange this “status quo” contract for a different contract. The intuition behind this result is that, under resolute choice, asymmetric information may differentially change the agents’ perception of the set of possible contracts. I provide necessary and sufficient restrictions on preferences for the no-trade result to hold under resolute choice models. The additional condition—*conditional decomposition*—is related to the union consistency property of decision procedures required for the general Agreement Theorem (Rubinstein and Wolinsky, 1990). Conditional decomposition together with dynamic consistency imply that the ex-ante preferences satisfy *weak decomposition*, which was characterized by Grant et al. (2000).

## 2. The trading environment

Assume a finite state space  $\Omega$ , with a generic element  $\omega$ .  $\Sigma$  is the algebra of events on  $\Omega$ . For every  $E \in \Sigma$  let  $E^c = \Omega \setminus E$  be the complement of  $E$ . Let  $I$  be a finite set of agents. The set of consequences is denoted by  $C$ . A contingent contract is a function from  $\Omega$  to  $C$ . Let  $A$  be the set of all contingent contracts. For every two contracts  $a, b \in A$  let the *composite contract*  $a$  at  $E$  given  $b$  be:

$$aEb = \begin{cases} a(\omega), & \omega \in E, \\ b(\omega), & \omega \in E^c. \end{cases} \quad (1)$$

Savage’s (1972) *Sure Thing Principle* asserts that preference between two contracts should not depend on those states in which the two contracts have the same consequences.

**The Sure Thing Principle P2.** For all non-null<sup>1</sup> events  $E$  and contracts  $a, b, h, h' \in A$ :  $aEh \succ^i bEh$  if and only if  $aEh' \succ^i bEh'$ .

This paper focuses on preferences that do not abide by the Sure Thing Principle. Uncertainty aversion is one example of such preferences. This paper does not specify a functional form for preferences however, thus allows for other possibilities. Following Machina (1989) *consequentialism* is relaxed in order to maintain dynamic consistency of preferences in non-expected utility models. The dynamic structure of preferences that results from this relaxation is referred to as *resolute choice*.

<sup>1</sup> An event  $E$  is null if and only if conditional on  $E$ , the agent is indifferent between any two contracts.

The decision process is divided into two non-identical stages. Wakker (1997) calls this *committed updating*.<sup>2</sup> Ex-ante, the decision maker chooses a contract, which serves later as a “status quo” contract. Once the decision maker learns that an event has not happened, her conditional preferences at the event which did obtain, depend on what would have happened at the non-realized event, *according to the original contract*. I.e., the contract which was a *real possibility* when those states were considered. The consistency condition in this framework implies that the agent’s preferences over contracts at the event which happened, conform to her preferences ex-ante, taking borne uncertainties (at the event which did not obtain) as fixed. Thus, when the decision maker learns that an event has occurred, she evaluates the contracts at that event in conjunction with the non-realized consequences of the status quo contract.

Each agent is endowed with a set of binary relations,  $\{\succ_E^i\}_{E \in \Sigma}$ , which represent agent  $i$ ’s conditional preferences over complete contracts at every event  $E$  (for ex-ante preferences the subscript  $\Omega$  is omitted). It is assumed that  $\succ^i$  is a weak order (Savage’s P1) and that the relation between ex-ante preferences and conditional preferences at events is governed by *dynamic consistency*, defined as follows.

**Definition 1.** Agent  $i$ ’s preferences over contingent contracts satisfy *Dynamic Consistency* if for all  $E, F \in \Sigma$  such that  $E \subseteq F$  and  $E$  is not  $\geq_F$ -null and for all contracts  $a, b, h \in A$ :  $(aEh \succ_E^i bEh)$  if and only if  $(aEh \succ_F^i bEh)$ .

Dynamic consistency implies that if two contracts differ only on  $E$ , the conditional preference between them at  $E$  should conform to the conditional preference between them when less information was available (and in particular, ex-ante, at  $\Omega$ ). *Consequentialism* implies that conditional preferences over contracts are independent of counterfactual outcomes.

**Definition 2.** Agent  $i$ ’s conditional preferences over contingent contracts satisfy *consequentialism* if for all non-null  $E \in \Sigma$  and all  $a, a', h, h' \in A$ :  $aEh \succ_E^i a'Eh$  if and only if  $aEh' \succ_E^i a'Eh'$ .

Conditional on an event, a decision maker acts as if she had started out from that point and treats the uncertainty that has not materialized as irrelevant or as if it never existed. Halevy (1998) shows that dynamic consistency and consequentialism imply the Sure Thing Principle, and that if preferences satisfy the Sure Thing Principle then the conditional preferences are consequentialist and dynamically consistent. Since I focus here on dynamically consistent preferences that are not separable (and hence do not satisfy the Sure Thing Principle), I concentrate on resolute (non-consequential) conditional preferences. For these preferences, the relation between the ex-ante and interim conditional preferences has to be detailed. In the following definition, a committed updating form of resolute preferences is adopted:

<sup>2</sup> Although committed updating is the prevalent non-consequential model that retains dynamic consistency, it is not the only one. Wakker (1997) proposes to consider *strategic updating* model.

**Definition 3.** Let  $b \in A$  be a status quo contract at a non-null event  $E \in \Sigma$ . Agent  $i$ 's *resolute conditional preferences over contracts at event  $E$  relative to contract  $b$*  (denoted by  $\succ_{E,b}^i$ ) are defined by: for any two contracts  $a, a' \in A$ ,  $a \succ_{E,b}^i a'$  if and only if  $aEb \succ_E^i a'Eb$ .

Definition 3 allows for conditional preferences between contracts at  $E$  to depend on the counterfactual consequences of the status quo contract  $b$ . It implies that a resolute decision maker evaluates a contract  $a \in A$  at an event  $E \in \Sigma$ , by considering the composite contract  $a$  at  $E$  given  $b$ , created by  $a$  in conjunction with the counterfactual part ( $E^c$ ) of the status quo contract. Using her conditional preferences, the decision maker evaluates the complete composite contract  $aEb$ . Note that the resolute conditional preferences defined above allows for consequentialism. The conditional preferences  $\succ_E^i$  defined on composite contracts of the sort:  $a$  at  $E$  given  $b$ , may be consequential (that is, independent of  $b$ ). In the absence of consequentialism, dynamic consistency and resolute choice (with committed updating) govern the relation between ex-ante and interim preferences.

I use a weak notion of ex-ante efficiency: 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  $i \in I$ .<sup>3</sup>

The informational environment is standard. Every agent is endowed with a knowledge function  $k^i$  on  $\Omega$  such that  $k^i(\omega)$  represents all the information agent  $i$  has at state  $\omega$ .<sup>4</sup> Let  $\Pi^i(\omega) = \{\omega' \in \Omega: k^i(\omega') = k^i(\omega)\}$  be the set of all states that are indistinguishable for agent  $i$  from state  $\omega$ . These classes of states (cells) constitute the information partition of the agent. An event  $H$  obtains at  $\omega$  if and only if  $\omega \in H$ . Agent  $i$  knows event  $H$  at  $\omega$  if  $\Pi^i(\omega) \subseteq H$ . The event “ $i$  knows  $H$  happened” is the set of states at which  $i$  knows  $H$  obtains:  $\{\omega: \Pi^i(\omega) \subseteq H\}$ . An event  $E$  is *self-evident* if for every  $i \in I$ :  $E = \{\omega: \Pi^i(\omega) \subseteq E\}$ . I.e., an event is self-evident if at the time of its happening all agents know it. An event  $B$  is *common knowledge* among  $I$  at  $\omega$  if there exists a self-evident event  $E$  such that  $\omega \in E$  and  $E \subseteq B$  (Aumann, 1976).

### 3. The possibility of speculative trade

The no-trade theorem states that if a status quo contract is ex-ante Pareto efficient, asymmetric information cannot lead rational agents to agree on a different contract that is commonly known to dominate the status quo contract. In other words, there exists no state in which the event that the other contract Pareto dominates (conditional on the asymmetric information) the ex-ante efficient contract is common knowledge.

One of the assumptions imbedded in the different no-trade theorems is expected utility, and especially its additive separability property. The latter is a result of Savage's Sure

<sup>3</sup> The stronger version of ex-ante efficiency: there does not exist a contract  $a$  such that  $a \succsim^i b$  for all agents and there exists at least one agent for whom  $a \succ^i b$ , would not change the argument.

<sup>4</sup> Alternatively, we could think of  $k^i(\omega)$  as a signal the agent receives when the state is  $\omega$ . The assumption that  $k^i$  is a function (and not a correspondence) is equivalent to assuming that the information (signal) is without noise.

Thing Principle. Dow et al. (1990) show that departures from expected utility in which consequentialism is retained, that is, the decision maker treats counterfactual consequences as irrelevant at the interim stage, makes speculative trade possible. Note that if separability is relaxed and consequentialism is retained, dynamic consistency of preferences is lost. Halevy (1998) shows that the trade result in this case does not rely on asymmetric information, and could be supported within a symmetric information environment. This section shows that dynamic consistency alone is not sufficient to preclude speculative trade. Particularly, by relaxing consequentialism, it is possible to generate speculative trade while maintaining dynamic consistency.

Assume there are two possible states of the world, so that  $\Omega = \{\omega_1, \omega_2\}$ , and two agents Alice (A) and Bob (B), who are dynamically consistent but non-consequentialist. The information structures are:

$$\Pi^A = \{(\omega_1), (\omega_2)\} \quad \text{and} \quad \Pi^B = \{(\omega_1, \omega_2)\}. \quad (2)$$

That is, when a state of nature has occurred, Alice knows it while Bob remains ignorant. The set of consequences  $C$  is:  $\{a_1, a_2, b_1, b_2\}$ , where  $a_1$  and  $b_1$  are possible consequences at state  $\omega_1$ , while  $a_2$  and  $b_2$  are possible consequences at  $\omega_2$ . I do not restrict the set of consequences, therefore they could be lotteries (dependent on aggregate uncertainty), allocation in a stochastic economy or any other structure on which preferences are defined. The set of contracts is:  $A = \{(a_1, a_2), (a_1, b_2), (b_1, a_2), (b_1, b_2)\}$ . Let Alice and Bob's preferences be:

$$\begin{aligned} (a_1, b_2) &\sim^A (b_1, a_2) \succ^A (b_1, b_2) \succ^A (a_1, a_2) \\ (a_1, b_2) &\sim^B (b_1, a_2) \prec^B (b_1, b_2) \prec^B (a_1, a_2). \end{aligned} \quad (3)$$

Assume the ex-ante Pareto efficient contract  $(b_1, b_2)$  is the *status quo* contract (note that all contracts are ex-ante Pareto efficient). At the interim stage, both agents will find it beneficial to recontract to  $(a_1, a_2)$ . Alice learns the true state of the world and she prefers  $(a_1, a_2)$  to  $(b_1, b_2)$  relative to  $(b_1, b_2)$  at every state of the world. Formally, for  $i = 1, 2$ :

$$(a_1, a_2) \succ_{\{\omega_i, (b_1, b_2)\}}^A (b_1, b_2) \quad (4)$$

if and only if (by Definition 3 of resolute conditional preferences)

$$(a_1, a_2)\{\omega_i\}(b_1, b_2) \succ_{\{\omega_i\}}^A (b_1, b_2) \quad (5)$$

if and only if (by Definition 1 of dynamic consistency)

$$(a_1, a_2)\{\omega_i\}(b_1, b_2) \succ^A (b_1, b_2). \quad (6)$$

Bob remains ignorant (i.e.,  $\Pi^B(\omega_i) = \Omega$  for  $i = 1, 2$ ). Since Alice would suggest that the contract  $(a_1, a_2)$  would be signed in every state, and he prefers  $(a_1, a_2)$  to  $(b_1, b_2)$ , Bob will find it profitable to agree to the trade. Formally, the event:

$$G = \{\omega: (a_1, a_2) \succ_{\Pi^i(\omega), (b_1, b_2)}^i (b_1, b_2) \text{ for all } i \in I\} = \{\omega_1, \omega_2\} = \Omega \quad (7)$$

is common knowledge at every state. That is, it is common knowledge at every state that all agents prefer to recontract to  $(a_1, a_2)$  from  $(b_1, b_2)$ . Hence, it is commonly known that there exist a contract that interim Pareto dominates an ex-ante efficient contract.

What drives the trade in this example is that Alice's preferences between  $(a_1, a_2)$  and  $(b_1, b_2)$  are reversed at the interim stage.<sup>5</sup> For her, the interim evaluation of  $(a_1, a_2)$  in state  $\omega_1$  reflects the fact that she receives  $a_1$  and foregoes  $b_2$ . As a result, in Alice's eyes, the individual consequences specified by the contract  $(a_1, a_2)$  are different at the ex-ante and the interim stage. Hence  $(a_1, a_2)$  is perceived as being two different contracts as the true state of the world is learned. However, from Bob's point of view, the consequences of the contracts do not change between the ex-ante and the interim stages, since he does not learn the true state of the world until the end of the trading game.

Note that if Bob had learned the true state of the world at the interim stage, there would not have been any trade: asymmetric information is crucial for this argument. Contrary to this, if consequentialism is retained and dynamic consistency is dropped (Dow et al., 1990) asymmetric information is not necessary to generate trade (Halevy, 1998, Appendix B). Hence, the pattern of trade presented here is compatible with Geanakoplos' (1992) terminology of *speculative trade*: ex-ante efficiency is not enough to prevent trade, because agents' perception of the set of contracts may change at the interim stage as a result of asymmetric information and the committed updating used in the resolute choice preferences.

The example proves that even in the simplest circumstances (two states and noiseless information) asymmetric information may lead rational agents, whose preferences do not conform to the Sure Thing Principle, to trade. The set of preferences that supports this example is not degenerate. In Halevy (1998) I show that non-neutrality of uncertainty (aversion or love) may produce similar examples. Other applications could rely on disappointment aversion (Gul, 1991) in which the relaxation of consequentialism is very natural.

#### 4. No-trade theorem for resolute choice

The above example relies on a violation of a principle required for decision procedures (defined below) in the general Agreement Theorem. This theorem is presented first, and from it I derive the necessary and sufficient conditions on preferences for a no-trade theorem in a noiseless environment.

##### 4.1. Preliminary: The Agreement Theorem

Let  $Z^i$  be agent  $i$ 's set of possible actions. A *decision procedure* of agent  $i$  is a function from  $\Sigma$  to  $Z^i$  such that  $D^i(E)$  is the recommendation of the decision procedure to agent  $i$  with information  $E$ . A decision procedure  $D^i$  satisfies *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 union consistency. An *action function*  $d^i$  is the action of agent  $i$  at state  $\omega$ , or the function that implements  $D^i$  at  $\omega$ .

<sup>5</sup> On the Dutch book potential in the non-consequentialist preferences, which is present in the above example, see Machina's (1989) discussion on Segal and Dekel's examples and Wakker's (1997) discussion on strategic updating.

Formally,  $d^i : \Omega \rightarrow Z^i$  such that  $d^i(\omega) = D^i(\Pi^i(\omega))$ . Define the event “ $i$  takes action  $z^i$ ” as:  $\{d^i = z^i\} \equiv \{\omega \in \Omega : d^i(\omega) = z^i\}$ .

**The Agreement Theorem** (Rubinstein and Wolinsky, 1990; Geanakoplos, 1992). *Let  $z^i \in Z^i$  for all  $i \in I$ . If  $\bigcap_{i \in I} \{d^i = z^i\}$  is common knowledge at  $\omega$ , 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  $\omega$  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  $E$  for every  $i \in I$ .

#### 4.2. Characterization of no-trade

This section presents necessary and sufficient conditions on preferences for the no-trade theorem to hold under a noiseless information structure. The first condition, *conditional decomposability*, is related to the notion of union consistency of decision procedures. The effect of this condition is to eliminate situations like the one in the example where, for  $i = 1, 2$ , Eq. (4) holds but  $(b_1, b_2) \succ^A (a_1, a_2)$ .

**Definition 4.** A preference structure  $\{\succ_E^i\}_{E \in \Sigma}$  over contingent contracts satisfies *conditional decomposability* if for all  $a, b \in A$ , all non-null  $E \in \Sigma$  and all partitions  $E_1, \dots, E_n$  of  $E$ : if  $a \succ_{E_j, b}^i b$  for every  $j = 1, \dots, n$  then  $a \succ_{E, b}^i b$ .

This property states that given a “status quo” 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. Note that the contract  $b$  is the resolute choice contract appearing in  $\succ_{E, b}^i$ . The definition of conditional decomposability does not imply that if  $a \succ_{E_j, b}^i c$  for all  $j$  then  $a \succ_{E, b}^i c$  for  $c \neq b$ . Conditional decomposability is weaker than consequentialism and applies to non-separable preferences. The following proposition proves that the above condition, together with dynamic consistency, is sufficient to imply the no-trade result.

**Proposition 1.** *Let all agents have resolute conditional preferences  $\{\succ_E^i\}_{E \in \Sigma}$  over contingent contracts that satisfy dynamic consistency and conditional decomposability. If  $b$  is an ex-ante efficient contract, then there exist no state  $\omega^*$  and another contract  $a$ , such that the event:*

$$G = \{\omega : a \succ_{\Pi^i(\omega), b}^i b \text{ for all } i \in I\} \tag{8}$$

*is common knowledge at  $\omega^*$ .*

*Furthermore, if either dynamic consistency or conditional decomposability are not satisfied there exist informational structures and a contract set such that  $b, a$  and  $\omega^*$  exist and  $G$  above is common knowledge at  $\omega^*$ .*

**Proof.** 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. Agent’s  $i$  decision procedure is

$$D^i(E, a, b) = \begin{cases} T, & a \succ_{E,b}^i b \\ N, & \text{otherwise.} \end{cases} \quad (9)$$

Note that  $D^i$  depends on the counterfactual consequences of the status quo contract  $b$ . Agents will trade  $a$  for  $b$  if and only if there exists a state  $\omega^*$  at which it is common knowledge that all agents are willing to trade. Since  $\{\succ_E^i\}_{E \in \Sigma}$  satisfies conditional decomposability,  $D^i$  satisfies union consistency. Assume such  $\omega^*$  and  $a$  exist. Therefore the event  $G$  which could be written as:  $G = \bigcap_{i \in I} \{D^i(E, a, b) = T\}$  is common knowledge at  $\omega^*$ . By the Agreement Theorem, there exists an event  $E$ , such that  $D^i(E, a, b) = T$ . That is:  $a \succ_{E,b}^i b$  for all agents. But then by dynamic consistency:  $aEb \succ^i b$  for all agents, which contradicts the assumption that  $b$  was ex-ante efficient.

Necessity is proved directly by construction.  $\square$

Proposition 1 is a contribution to the existing literature of no-trade theorems. Previous works (e.g., Rubinstein and Wolinsky, 1990) which started either from expected utility or general decision functions, did not distinguish between dynamic consistency and conditional decomposability, a distinction which exists for any preference relation that does not satisfy the Sure Thing Principle. Thus, Proposition 1 specifies the conditions which those general preferences have to satisfy in order for the no-trade result to hold.

## 5. Concluding remarks

(1) Proposition 1 could be interpreted as follows: If  $b$  ex-ante dominates  $a$  for all agents then there exists no state at which it is common knowledge that  $a$  dominates  $b$  for all agents given their information. Hence it gives sufficient and necessary conditions on preferences for consistency of the ex-ante and the interim Pareto ranking.

(2) Proposition 1 includes expected utility as a special case. Definition 3 of resolute conditional preferences includes consequential conditional preferences as a special case. As mentioned earlier, dynamic consistency and consequentialism imply the Sure Thing Principle, which together with Savage’s P1 imply conditional decomposability (Halevy, 1998). Hence, the proposition includes as a special case the results of Rubinstein and Wolinsky (1990) on no-trade, who employ a similar methodology but consider expected utility and a common prior, and the result of Dow et al. (1990) who prove the necessity of additivity for the no-trade result when considering only consequential preferences. In this set, only preferences that abide by the Sure Thing Principle will be dynamically consistent.

(3) The two conditions in Proposition 1, conditional decomposability and dynamic consistency, imply Grant’s et al. (2000) weak decomposability:



A preference relation  $\succ^i$  over contingent contracts satisfies *weak decomposability* if for all  $a, b \in A$ , and events  $E \subseteq \Omega$  such that  $aEb \succ^i b$  and  $bEa \succ^i b$  then  $a \succ^i b$ .

Grant et al. give an important characterization of this property that enables the identification of preference relations where speculative trade would be possible. They prove that for probabilistically sophisticated preferences (Machina and Schmeidler, 1992; Epstein and Le Breton, 1993), weak decomposability is equivalent to the Betweenness property, and give an implicit additive representation of preferences in the absence of probabilistic sophistication. This representation intersects (but does not overlap) ambiguity averse preferences. Grant et al. relate weak decomposability to Gul and Lantto's (1990) "Dynamic Programming Solvability" property, for choice among and along (objective) decision trees under resolute choice. Hence resolute choice and dynamic consistency relate weak decomposability of the ex-ante preference relation to conditional decomposability of the interim preferences.

### Acknowledgments

I am grateful to a referee and an associate editor of this journal for very valuable remarks. I benefited from discussions on the subject with Luca Anderlini, Robert Aumann, Tomer Blumkin, Yossi Feinberg, Sergiu Hart, Atsushi Kajii, Zvika Neeman, George Mailath, Jacob Sagi, Oscar Volij, Peter Wakker, and Asher Wolinsky. Remarks from seminar participants at the Hebrew University, Ben Gurion University, University of Pennsylvania, Brown University, and University of British Columbia are gratefully acknowledged.

This work is based on part of the first chapter of my PhD dissertation, written under the supervision of Professor Menahem Yaari at the Hebrew University of Jerusalem. I thank him for his advice and support. I thank the Department of Economics at the University of Pennsylvania for its hospitality while working on this paper.

### References

- Aumann, R.J., 1976. Agreeing to disagree. *Ann. Statist.* 4, 1236–1239.
- Dow, J., Madrigal, V., Werlang, S.R.C., 1990. Preferences, common knowledge, and speculative trade. Working Paper IFA-125-90. London Business School.
- Epstein, L., Le Breton, M., 1993. Dynamically consistent belief must be Bayesian. *J. Econ. Theory* 61, 1–22.
- Geanakoplos, J., 1992. Common knowledge. *J. Econ. Perspect.* 6 (4), 53–82.
- Grant, S., Kajii, A., Polak, B., 2000. Decomposable choice under uncertainty. *J. Econ. Theory* 92 (2), 169–197.
- Gul, F., 1991. A theory of disappointment aversion. *Econometrica* 59 (3), 667–686.
- Gul, F., Lantto, O., 1990. Betweenness satisfying preferences and dynamic choice. *J. Econ. Theory* 52, 162–177.
- Halevy, Y., 1998. Trade between rational agents as a result of asymmetric information. CARESS Working Paper 98-03. University of Pennsylvania.
- Hammond, P.J., 1988. Consequentialist foundation for expected utility. *Theory Dec.* 25, 25–78.
- Holmström, B., Myerson, R.B., 1983. Efficient and durable decision rules with incomplete information. *Econometrica* 51, 1799–1820.
- Machina, M.J., 1989. Dynamic consistency and non-expected utility models of choice under uncertainty. *J. Econ. Lit.* 27, 1622–1668.

- Machina, M.J., Schmeidler, D., 1992. A more robust definition of subjective probability. *Econometrica* 60 (4), 745–780.
- Milgrom, P., Stokey, N., 1982. Information, trade and common knowledge. *J. Econ. Theory* 26, 17–27.
- Rubinstein, A., Wolinsky, A., 1990. On the logic of ‘agreeing to disagree’ type results. *J. Econ. Theory* 51, 184–193.
- Savage, L.J., 1972. *The Foundations of Statistics*, 2nd revised edition. Dover, New York.
- Wakker, P., 1997. Are counterfactual decisions relevant for dynamically consistent updating under nonexpected utility. CentER Working Paper.