

Dynamic Inconsistency and Inefficiency of Equilibrium under Knightian Uncertainty*

Patrick Beissner[†] Michael Zierhut[‡]

March 14, 2019

Abstract

This paper extends the theory of general equilibrium with Knightian uncertainty to economies with more than two dates. Agents have incomplete preferences with multiple priors à la Bewley. These priors are updated in light of new information. Contrary to the two-date model, the market outcome varies with choice of updating rule. We document two phenomena: First, agents may find it optimal to deviate from their initial trading plans. This dynamic inconsistency may result in Pareto inefficiency, even if the market is dynamically complete. Second, ambiguous probability mass may spread; new information may create new uncertainty. Either phenomenon is avoided under certain updating rules: Full Bayesian updating guarantees dynamic consistency and Pareto efficiency, while maximum-likelihood updating prevents ambiguity spillovers. We ask whether it is possible to design updating rules that combine both properties. The answer is negative: No such rule exists.

JEL Classification: D52; D53; D81

Keywords: Knightian uncertainty, Dynamic consistency, General equilibrium, Incomplete markets

*The authors thank Pablo Beker, Federico Echenique, Georgios Gerasimou, Christian Kellner, Fabio Maccheroni, Filipe Martins-da-Rocha, Frank Riedel, Ronald Stauber, Georg Weizsäcker, and Jan Werner, as well as participants in the 2018 North American Summer Meeting of the Econometric Society and the workshop on Mathematics of Behavioral Economics and Knightian Uncertainty in Financial Markets at ZiF Bielefeld for helpful comments. M. Zierhut gratefully acknowledges financial support by the German Research Foundation (DFG) under project ZI 1673/1-1.

[†]Research School of Economics, ANU, 26012 Canberra, Australia, patrick.beissner@anu.edu.au

[‡]Institute of Financial Economics, Humboldt University, 10099 Berlin, Germany, michael.zierhut@hu-berlin.de.

1 Introduction

Two of the central themes of economic theory are decisions and markets. Consider the sequential decision problem of a market participant as information unfolds gradually over time. There is little disagreement among economists that better information should reduce uncertainty, but substantially more disagreement about the exact meaning of uncertainty. In expected utility theory, uncertainty is represented by states of the world, and agents assign probabilities – this is uncertainty in the traditional sense. By contrast, Knight (1921) reserves the term uncertainty for circumstances that do not permit the assignment of precise probabilities – uncertainty in the Knightian sense. Both types of uncertainty are synthesized in the Knightian decision model of Bewley (1986, 2002), which is a multiple-prior generalization of the subjective expected utility model of Savage (1954). Some events have ambiguous probabilities, but the agent has no attitude toward ambiguity: Knightian uncertainty results in incomplete preferences. The purpose of the present paper is to study competitive equilibria when agents with Bewley-type preferences meet in a multi-period financial market.

We build on the important contribution of Rigotti and Shannon (2005), who generalize the model of Arrow and Debreu (1954) to a setting with multiple priors. Two of their findings are particularly relevant in our context: First, Knightian uncertainty results in indeterminacy in equilibrium prices and allocations. Second, all equilibria are Pareto efficient. This welfare result is strong but limited to a very particular type of market: The only traded assets are state-contingent claims, and trade takes place only once. After the market has closed, the state of the world is revealed and all agents consume. We call this the two-date model, but it may also accommodate contingent claims for multiple future dates, as in Dana and Riedel (2013). The present paper considers a more general financial market, in which there are long-lived assets with repeated dividend payments, and these are traded at several consecutive dates. Thus, we generalize the two-date model with Knightian uncertainty in the same way as Radner (1972) generalizes the traditional Arrow-Debreu model.

The fundamental difference from the two-date model is that agents receive new information on several consecutive dates. Their information processing is represented by an updating rule that translates prior probabilities into posterior probabilities. In the traditional single-prior setting, Bayes' law determines the unique updating rule. In a setting with multiple priors, this is no longer true: A large set of updating rules is consistent with Bayes' law, and different agents may apply different rules. This matters in an equilibrium of plans, prices, and price expectations: Even though all agents form common and consistent price expectations and plan their trades and consumption in advance, some may find it optimal to deviate from their plans once they receive better information. These agents conclude that their plans were made under priors that are no longer consistent with the observed data. This leads to dynamically inconsistent decisions. We characterize dynamic consistency at the equilibrium level and find that dynamically inconsistent equilibria can be pervasive:

1. There are economies in which all equilibria are dynamically consistent and economies

in which all equilibria are dynamically inconsistent (Theorem 1 and Corollary 1).

As a consequence, dynamic consistency is not suited as a criterion of equilibrium refinement because it may eliminate all equilibria, or none. This raises the following question: Are dynamically inconsistent equilibria problematic from a normative viewpoint? To address this question, we apply two normative criteria. The first is Pareto efficiency, the most common criterion of social optimality. The second is that agents do not forget previous knowledge, which is a fundamental criterion of individual rationality.

As regards social optimality, it must be noted that financial market equilibria are typically not Pareto efficient when the financial market is incomplete. To eliminate confounding welfare effects of market imperfections, we study the special case of a dynamically complete market. Classical results of Kreps (1982) and Magill and Shafer (1990) show that the payoff space of a complete contingent claims market is spanned by a dynamically complete financial market. Even in such a perfect market, dynamic inconsistency may result in Pareto inefficient outcomes. We construct a tractable example in which all financial market equilibria are Pareto inefficient, in spite of dynamic completeness. Such undesirable market outcomes can be eliminated if the choice of updating rules is restricted. Dynamic consistency and efficiency are restored if all agents apply the full Bayesian rule:

2. Full Bayesian updating guarantees dynamic consistency and Pareto efficiency of equilibria with a dynamically complete market (Corollary 2 and Corollary 3).

As regards individual rationality, we require that agents use information to reduce uncertainty. Put differently, a release of new information must not create new uncertainty. However, it is easy to construct examples where this criterion is violated: Under some updating rules, ambiguous probability mass spreads to unambiguous events. Agents become uncertain about probabilities they knew before. We define a condition of ambiguity containment that prevents such ambiguity spillovers. This condition is met if all agents apply the maximum-likelihood rule:¹

3. Maximum-likelihood updating guarantees ambiguity containment (Proposition 1).

It should be noted that the maximum-likelihood rule creates a propensity to trade after an information release. This stands in contrast to the phenomenon of market breakdowns documented in various multi-prior settings, first by Mukerji and Tallon (2001) in a replication economy with ambiguity-averse traders. In an economy with preferences of the Bewley type, this phenomenon takes the form of a no-trade result: Rigotti and Shannon (2005) show that trade in the contingent claims market is no longer individually rational when uncertainty grows large. For a financial market that opens sequentially, we find the contrary: No-trade equilibria do not exist as long as some agents apply the maximum-likelihood rule.

¹For an analysis and comparison of the full Bayesian rule and the maximum-likelihood rule see Gilboa and Schmeidler (1993).

Since both of our normative criteria are standard in the single-prior setting, a natural objective is to identify an updating rule that guarantees these jointly, even if agents have multiple priors. However, this objective cannot be achieved:

4. There exists no updating rule that satisfies Bayes' law and guarantees both dynamic consistency and ambiguity containment (Corollary 5).

In light of this impossibility result, we ask whether there is a nontrivial subset of economies in which this conflict between normative criteria disappears. This question is answered in the affirmative: We characterize unproblematic economies, in which full Bayesian updates and maximum-likelihood updates coincide. Our characterization is based on a simple criterion for information partitions: If some observable event and some unambiguous event are neither disjoint nor related by set inclusion, the two normative criteria are at odds; otherwise, the conflict disappears. This criterion is strictly weaker than the rectangularity condition of Epstein and Schneider (2003), which characterizes dynamic consistency in the maxmin expected utility model of Gilboa and Schmeidler (1989). The latter shares with the Knightian decision model of Bewley (1986, 2002) its representation of multiple priors but employs a decision criterion that is based on ambiguity aversion.²

Two relations between the Knightian decision model and maxmin expected utility are highlighted by our results. First, while Hanany and Kilbanoff (2007) show that there exists no updating rule that satisfies Bayes' law and dynamic consistency under maxmin expected utility, we prove that such a rule exists in our setting: It is the well-known full Bayesian rule. Thus, dynamic consistency is a stronger requirement under maxmin expected utility. Second, if one resorts to rectangularity to ensure dynamic consistency, the conflict between normative objectives goes unnoticed. The rectangularity condition is sufficiently strong that all equilibria are dynamically consistent and free of ambiguity spillovers at the same time. This is true for both decision theoretic models, although the quantification *all equilibria* is weaker under maxmin expected utility. For the latter, the results of Rigotti and Shannon (2012) show that equilibria are generically determinate. Interestingly, there are exceptions: Mandler (2013) points out that indeterminacy can be robust if agents have access to a production technology.

The remainder of this paper is structured as follows. In Section 2, the model is introduced. In Section 3, equilibria under different updating rules are studied in a tractable example. In Section 4, all general results are derived and discussed. Section 5 concludes.

2 Model

Consider a stochastic finance economy with a finite state space Ω . The economy is populated by a finite number I of agents, who plan their consumption at the present date 0 and over a finite number T of future dates.

²Gilboa, Maccheroni, Marinacci, and Schmeidler (2010) construct a bridge between these two models.

2.1 Uncertainty and Knowledge

At date 0, the state of the world is drawn from Ω but not revealed to the agents. Each agent i assigns subjective probabilities $\pi^i(\omega)$ to all states $\omega \in \Omega$. The agent may have multiple priors, each represented by a different probability vector $\pi^i \in \Delta^{|\Omega|}$, which is the unit simplex in $\mathbb{R}^{|\Omega|}$. All these probability vectors are collected in a set Π^i , which satisfies the following assumption:

Assumption 1 (Subjective probabilities). For each agent i , Π^i is closed and convex. For each state $\omega \in \Omega$, there is some $\pi^i \in \Pi^i$ such that $\pi^i(\omega) > 0$.

At each date $t > 0$, new information about the state of the world is released. This gradual revelation is modeled as a sequence of information partitions $\mathcal{I} = \{\mathcal{I}_t\}_{t=0}^T$. Each event $\xi \in \mathcal{I}_t$ is an information set: It consists of all those states that cannot be distinguished on the basis of the information available at date t . The *information structure* \mathcal{I} represents uncertainty in the traditional sense. Initially, at date 0, none of the states can be distinguished; that is, $\mathcal{I}_0 = \{\Omega\}$. The information partitions become finer as time progresses. Finally, at date T , the state of the world is known; that is, $\mathcal{I}_T = \{\{\omega\}_{\omega \in \Omega}\}$. Thus, the information structure defines a tree of *date-events* (t, ξ) .

Each agent i updates his subjective probabilities on observing an event $\xi \in \Omega$. The *updating rule* of agent i is a closed-convex-valued correspondence $\Pi^i : 2^\Omega \rightrightarrows \Delta^{|\Omega|}$ from the observed event ξ to his set of posterior probabilities $\Pi^i(\xi)$. This rule must satisfy $\Pi^i(\Omega) = \Pi^i$, and since we follow the convention of dropping parentheses whenever Ω is the argument, there is no notational ambiguity between the updating rule and the set of prior probabilities. Bayesian updating is the natural rule in the single-prior case: The unique posteriors are the conditional probabilities

$$\pi^i(\omega|\xi) = \begin{cases} \frac{\pi^i(\omega)}{\pi^i(\xi)} & \text{if } \omega \in \xi \\ 0 & \text{otherwise} \end{cases} . \quad (1)$$

By contrast, in the multi-prior case, the set of candidate updating rules is large, and the literature has not yet reached a consensus. A minimal requirement is Bayes' law:

Definition 1. An updating rule Π^i satisfies *Bayes' law* if for any $\xi \subseteq \Omega$ the following condition is true:

$$\pi^i \in \Pi^i(\xi) \implies \exists \hat{\pi}^i \in \Pi^i \text{ with } \pi^i(\cdot) = \hat{\pi}^i(\cdot|\xi) .$$

Each agent may have his individual updating rule, as long as this rule is consistent with Definition 1. As we shall see, two candidates for Π^i are particularly relevant from a normative viewpoint. The first is the so-called *full Bayesian rule*:

$$\Pi_B^i(\xi) = \{ \pi^i(\cdot|\xi) \mid \pi^i \in \Pi^i \} . \quad (2)$$

It applies Equation (1) prior by prior and is therefore the least selective updating rule that satisfies Bayes' law. By contrast, only those priors that assign the highest probability to the observed event are considered under the *maximum-likelihood rule*:

$$\Pi_{\mathbb{M}}^i(\xi) = \left\{ \pi^{i*}(\cdot | \xi) \mid \pi^{i*} \in \arg \max_{\pi^i \in \Pi^i} \pi^i(\xi) \right\}. \quad (3)$$

In general, the updating rule determines how the set of subjective probabilities evolves over time and is thus a complete description of Knightian uncertainty as faced by agent i . If it assigns a unique probability to an event, this event is called *unambiguous*; if it assigns multiple probabilities, the event is called *ambiguous*. Since uncertainty depends on information, the classification of unambiguous and ambiguous events may change over time. This classification can be represented as a sequence of ambiguity partitions. Let pr_{ζ} denote the projection onto the subspace of those states contained in ζ .

Definition 2. The *ambiguity partition* $\mathcal{A}^i(\xi)$ of agent i , conditional on event ξ , is the finest possible partition of Ω such that the set of posterior probabilities can be decomposed as the Minkowski sum

$$\Pi^i(\xi) = \sum_{\zeta \in \mathcal{A}^i(\xi)} \text{pr}_{\zeta}(\Pi^i(\xi)). \quad (4)$$

In other words, after ξ is observed, $\mathcal{A}^i(\xi)$ is the finest partition that contains only unambiguous events. There is no uncertainty about the probabilities of its cells: If $\zeta \in \mathcal{A}^i(\xi)$, then $\pi^i(\zeta) = \hat{\pi}^i(\zeta)$ for any $\pi^i, \hat{\pi}^i \in \Pi^i(\xi)$. All events in $\mathcal{A}^i(\Omega) = \mathcal{A}^i$ are unconditionally ambiguous.³

The *knowledge* of agent i at any date-event (t, ξ) , with $\xi \in \mathcal{I}_t$, is summarized by the tuple $(\xi, \Pi^i(\xi))$, which consists of his information set and his set of posterior probabilities. The normative criterion that *agents do not forget what they knew*, requires that information partitions and ambiguity partitions become progressively finer over time: If $t > s$, then \mathcal{I}_t is finer than \mathcal{I}_s ; if $\xi \subseteq \xi'$, then $\mathcal{A}^i(\xi)$ is finer than $\mathcal{A}^i(\xi')$. If one of these two conditions is violated, new information creates new uncertainty, either in the traditional or Knightian sense.

2.2 Preferences and Decisions

Let \mathcal{S} denote the space of all real-valued stochastic processes that are measurable with respect to the filtration generated by \mathcal{I} , with the inner product $x \cdot y = \sum_{t=1}^T \sum_{\xi \in \mathcal{I}_t} x_t(\xi) y_t(\xi)$. By contrast, $x \otimes y$ denotes the componentwise product. The consumption set $\mathcal{S}_+ = \{x \in \mathcal{S} \mid x \geq 0\}$ is defined as the cone of nonnegative processes. Note that $x \geq 0$ means all components are nonnegative, $x > 0$ means at least one component is not zero, and $x \gg 0$ means all components are greater than zero.

Each agent i has a preference relation $x \succ^i y$ over consumption plans $x, y \in \mathcal{S}_+$; that is, agent i prefers x to y . Let $\mathbb{E}^{\pi^i}[\cdot]$ denote the expectation operator under π^i . The preferences of all consumers are assumed to be of the following form:

Assumption 2 (Preferences). For each agent i ,

$$x \succ^i y \quad \text{if and only if} \quad \mathbb{E}^{\pi^i}[u^i(x)] > \mathbb{E}^{\pi^i}[u^i(y)] \quad \forall \pi^i \in \Pi^i,$$

³Epstein (1999) calls the members of \mathcal{A}^i and all their unions the *collection of unambiguous events*.

in which $u^i : \mathbb{R}_+^{T+1} \rightarrow \mathbb{R}$ is continuous, strictly increasing, strictly concave, and differentiable in \mathbb{R}_{++}^{T+1} .

Under Assumption 2, each \succ^i is completely described by the pair (Π^i, u^i) . If Π^i is a singleton, the agent is a traditional expected utility maximizer. However, if Π^i is a larger set, the agent's preferences are transitive, monotone, and convex but incomplete. Some plans are always comparable: If $x \gg y$, then the expected utility of x must exceed the expected utility of y under any possible probability distribution.

While the above formulation makes expected utility theory appear to be a special case of Knightian decision theory, it should be stressed that the decision criterion is fundamentally different: Expected utility maximizers base their decisions on a criterion of optimality. A decision is optimal if it is preferred to all its alternatives. By contrast, Knightian decision makers rely on a criterion of undominatedness. A decision is undominated if there is no preferred alternative.⁴ Since optimality implies undominatedness, the latter is a weaker criterion. In fact, it is a criterion too weak to rule out erratic behavior in the case of sequential decisions: Agents may switch back and forth between two undominated plans for no reason. To avoid such intransitive choices, Bewley (1986, 2002) introduces the behavioral assumption of *inertia*: An agent deviates from the status quo only if the status quo is dominated by the new consumption plan.

The ranking of plans may change as new information becomes available. The conditional preference relation of agent i on observing ξ is jointly determined by Assumption 2 and the update $\Pi^i(\xi)$:

$$x \succ_{\xi}^i y \quad \text{if and only if} \quad \mathbb{E}^{\pi^i}[u^i(x)] > \mathbb{E}^{\pi^i}[u^i(y)] \quad \forall \pi^i \in \Pi^i(\xi).$$

We shall write $x \succ_t^i y$ whenever $x \succ_{\xi}^i y$ for all $\xi \in \mathcal{I}_t$. Therefore, an agent's ranking of alternative consumption plans depends on his information, his updating rule, and his utility function.

2.3 Financial Market Equilibrium

Each agent i has a stochastic endowment $e^i \in \mathcal{S}_+$. Since the number of date-events is finite, any process in \mathcal{S} can be viewed as a vector whose components are indexed as follows: $e_t^i(\omega)$ represents the endowment at date t in state ω , and $e^i(\omega)$ is the $(T+1)$ -dimensional subvector that represents the endowment at all dates in state ω . If $\xi \in \mathcal{I}_t$, then $e_t^i(\xi)$ stands for the identical endowment in all states $\omega \in \xi$. The superscript is omitted whenever vectors are joined; that is, $e = (e^1, \dots, e^I) \in \mathcal{S}_+^I$.

Assumption 3 (Endowments). For each agent i , the endowment satisfies $e^i \gg 0$.

A stochastic economy is completely described by the information structure \mathcal{I} , the updating rules $\Pi = \Pi^1 \times \dots \times \Pi^I$, and the utility functions and endowments of all agents.

⁴Two decisions can both be undominated for two reasons: Either they are not comparable or the agent is indifferent. However, indifference can be a narrow concept in the Knightian decision model. Gerasimou (2018) shows that indifference sets are singletons if Π^i is full-dimensional.

Definition 3. A *stochastic economy* is a tuple (\mathcal{I}, Π, u, e) of information structure, updating rules, utility functions, and endowments.

To enable transfers of income between different date-events, the economy is equipped with a finite number J of long-lived assets. At each date $t > 0$, asset j pays a dividend $A_t^j(\omega)$ to its holder in state ω . After these dividend payments, the financial market opens and assets can be traded. The asset structure is completely described by the dividend process $A \in \mathcal{S}^J$, which is required to have identical payments in states that cannot be distinguished. In addition, a mild regularity assumption is made:

Assumption 4 (Asset structure). At each date $t > 0$, $A_t \geq 0$ and initially $A_0 = 0$.

Under Assumption 4, asset prices are never negative because dividends are never negative. Moreover, no dividend payments occur before the market opens for the first time. Therefore, it is safe to assume that all assets are in zero net supply and that no agent is born with an asset portfolio. Agents face a sequence of decisions: At date 0, each agent i makes a consumption plan $C_0^i \in \mathcal{S}_+$ and a trading plan $\Psi_0^i \in \mathcal{S}^J$. Set $C_0^i = (c_0^i, \dots, c_T^i)$ and $\Psi_0^i = (\psi_0^i, \dots, \psi_T^i)$. By convention, time subscripts for capital letters refer to the date when plans are *made*, while time subscripts for small letters refer to the date when plans are *executed*. These plans involve consumption c_0^i and portfolio ψ_0^i purchased in the date 0 market and also state-contingent consumption and portfolios $c_t^i(\omega)$ and $\psi_t^i(\omega)$ at all future dates $t > 0$. Thus, agents plan all future actions in advance, on the basis of the information available at date 0. However, at each consecutive date, new information becomes available, and some agents may want to adjust their plans. At any interim date s , that is, $0 < s \leq T$, agent i may switch to a conditionally preferred consumption and trading plan for all future dates $s \leq t \leq T$. Past trades cannot be undone.

This sequence of decisions results in a sequence of consumption plans $C^i = (C_0^i, \dots, C_T^i) \in \mathcal{S}_+^{T+1}$ and a sequence of trading plans $\Psi^i = (\Psi_0^i, \dots, \Psi_T^i) \in \mathcal{S}^{J(T+1)}$; one for each date. Note that capital letters are used to denote sequences of plans; in this case, subscripts denote the date a plan was made. Decisions are called dynamically consistent if the agent makes one consumption and trading plan at date 0 with no adjustments at later points in time: Even in the light of new information, the agent finds it optimal to stick to his original plans.

Definition 4. The plans (C^i, Ψ^i) of agent i are *dynamically consistent* if for each date $t > 0$, $C_t^i = C_{t-1}^i$ and $\Psi_t^i = \Psi_{t-1}^i$.

The decisions of each agent are constrained by a sequence of budget sets, which depend on asset prices. Let $p \in \mathcal{S}_+^J$ be the asset price process. Since no dividends are paid after date T , it is clear that $p_T = 0$, and there is no more trade at the terminal date. The portfolio payoff at date t in state ω ,

$$X_t(p, \psi^i; \omega) = A_t(\omega) \cdot \psi_t^i(\omega) + p_t(\omega) \cdot \left(\psi_{t-1}^i(\omega) - \psi_t^i(\omega) \right),$$

is the sum of dividend payments and the change in investment. It shall be understood that $X(p, \psi^i, \omega) = (X_0(p, \psi^i; \omega), \dots, X_T(p, \psi^i; \omega))^\top$, and $\psi_{-1}^i = 0$ by definition. If the trading plan carried over from date $t - 1$, that is, component $t - 1$ of Ψ^i , is denoted by $\hat{\psi}^i = \Psi_{t-1}^i$, the *financial market budget set* at date t is defined as:

$$B_t^i(p, \hat{\psi}^i) = \left\{ (c^i, \psi^i) \in \mathcal{S}_+ \times \mathcal{S}^J \mid \begin{array}{l} c^i(\omega) - e^i(\omega) = X(p, \psi^i; \omega) \quad \forall \omega \in \Omega \\ \psi_s^i = \hat{\psi}_s^i \quad \forall s < t \end{array} \right\}.$$

The sequence of decisions leads to a sequence of first-order conditions: At any date-event (s, ξ') , an interior plan is undominated if and only if for any consecutive date-event (t, ξ) (i.e., $s \leq t$ and $\xi \subseteq \xi'$),

$$p_t(\xi) = \mathbb{E}^{\pi^i} \left[\frac{\partial_{c_{t+1}^i} u^i[c^i]}{\partial_{c_t^i} u^i[c^i]} (A_{t+1} + p_{t+1}) \right] \quad \text{for some } \pi^i \in \Pi^i(\xi'). \quad (5)$$

Note that in Equation (5), information enters the decision problem through the updating rule $\Pi^i(\xi)$. At each date t , current asset prices p_t are determined in the financial market, and all agents act as price takers. By contrast, future asset prices are not yet known, and agents need to base their decisions on price expectations. As in Radner (1970, 1972) and Lucas and Prescott (1971), price expectations are part of the equilibrium concept. It is assumed that all agents have *common* expectations: On the basis of the available information \mathcal{I}_t , all agents form the same price expectations $P_t \in \mathcal{S}^J$. Thus, they associate the same prices with the same event, although they may disagree about its probability. Moreover, these expectations are *consistent* with their plans: Given P_t , their consumption and trading plans (C_t, Ψ_t) are such that supply equals demand in each of the consecutive markets. The common and consistent expectations of future prices may be revised at any date t as new information becomes available. Such revisions only take place if necessary; that is, if the old price expectations are no longer consistent with the plans of the agents.

This leads to a sequence of price expectations $P = (P_0, \dots, P_T) \in \mathcal{S}^{J(T+1)}$, in which P_T is the process of realized prices. If these prices are correctly anticipated at date 0 and $P_0 = \dots = P_T$, we speak of *correct expectations*. The concept of correct expectations is widespread under a variety of synonyms that include *rational expectations* and *self-fulfilling expectations*. It is closely related to dynamic consistency: When all agents make dynamically consistent decisions, their plans are never adjusted to new information, and the original price expectations are indeed self-fulfilling. Consistent expectations are a posteriori correct, and the system is closed. This automatism is embedded in Radner's definition of an equilibrium of plans, prices, and price expectations, which reduces the entire sequence (P, C, Ψ) to a single tuple (p, c, ψ) . We adopt Radner's concept of equilibrium, albeit without the preconceptions of dynamic consistency and correct expectations.

Definition 5. A *financial market equilibrium* for the stochastic economy (\mathcal{I}, Π, u, e) with asset structure A is a tuple (P, C, Ψ) of price expectations, consumption plans, and trading plans that satisfies at each date t ,

1. undominated choice: for each agent i ,

$$\nexists (\hat{c}^i, \hat{\psi}^i) \in B_t^i(P_t, \Psi_{t-1}^i) \text{ such that } \hat{c}^i \succ_t^i C_t^i$$

2. inertia: for each agent i ,

$$(C_t^i, \Psi_t^i) \neq (C_{t-1}^i, \Psi_{t-1}^i) \text{ only if } C_t^i \succ_t^i C_{t-1}^i$$

3. market clearing,

$$\sum_{i=1}^I \Psi_t^i = 0.$$

If all agents make dynamically consistent plans in a financial market equilibrium, we speak of a dynamically consistent equilibrium; otherwise, the equilibrium is called *dynamically inconsistent*. The behavioral assumption of inertia ensures that agents do not change their plans without good reason: If plans from the previous date $t - 1$ also satisfy the first-order conditions at date t , agents stick to these plans.⁵ These conditions are easier to verify when written in state price form. Recall that each asset price process p induces a set of state price processes

$$Q[p] = \{q \in \mathcal{S} \mid q \cdot X(p, \mathcal{S}^J) = 0\},$$

which are orthogonal to the marketed subspace. Equation (5) relates this set to the *marginal rates cone* of consumer i , defined as

$$\nabla U^i[c^i](\xi) = \{\alpha (\pi^i \otimes Du^i[c^i]) \mid \pi^i \in \Pi^i(\xi), \alpha \in \mathbb{R}_+\}, \quad (6)$$

which is the closed, convex cone that contains the marginal rates of substitution vectors under all posterior probabilities. $Du^i[c^i]$ is the gradient of u^i at c^i . Only if $\Pi^i(\xi)$ is a singleton is this cone a ray; in all other cases, the marginal rates cone has a nonempty relative interior. The first-order condition (5) at date 0 implies that⁶

$$Q[p] \cap \nabla U^i[c^i] \neq \emptyset. \quad (7)$$

At any date $t > 0$, this intersection becomes conditional on the observed information set ξ . Using the conditional state price set $Q[p](\xi) = \text{pr}_\xi(Q[p])$, the first-order condition is written as

$$Q[p](\xi) \cap \nabla U^i[c^i](\xi) \neq \emptyset. \quad (8)$$

⁵Rigotti and Shannon (2005) and Dana and Riedel (2013) use inertia with respect to agent endowments as a refinement that eliminates equilibria with individually irrational decisions. This refinement can be integrated into Definition 5 by setting $(C_{-1}^i, \Psi_{-1}^i) = (e^i, 0)$.

⁶The arguments that lead from (5) to (7) are standard and can be found, for example, in Magill and Quinzii (1996), Chapter 21.

Note that $Q[p]$ is a subspace of \mathcal{S} , whose dimension depends on how incomplete the financial market is. The market is called dynamically complete if at each date-event (t, ξ) , all direct successors in the date-event tree are be spanned by feasible portfolio payoffs. In this case, independent transfers of consumption to each successor are feasible. Stated rigorously:

Definition 6. The financial market is *dynamically complete* at (P, C, Ψ) if

$$\dim(X_{t+1}(P_t, \mathcal{S}^J; \xi)) = |\{\xi' \in \mathcal{I}_{t+1} \mid \xi' \subseteq \xi\}| \quad \forall \xi \in \mathcal{I}_t \quad \forall t < T.$$

Dynamic completeness must be defined relative to a given financial market equilibrium: Since all assets are long-lived, portfolio payoffs $X_{t+1}(P_t, \psi^i; \xi)$ depend on asset prices at date $t + 1$. If the financial market is dynamically complete and $p = P_0 = \dots = P_T$, then $Q[P_0]$ is one-dimensional and the first-order conditions (7) and (8) become simple set inclusions.⁷

It should be noted that financial market equilibria are at best determinate if Π^i is a singleton for every agent i . In the case of multiple priors, there is a continuum of equilibria. This continuum may involve exotic constellations such as equilibria with different allocations but the same prices, as well as equilibria with different prices but the same allocations.⁸

2.4 Contingent Market Equilibrium

In our welfare analysis, we make use of the results of Rigotti and Shannon (2005), who characterize Pareto efficiency under Knightian uncertainty and introduce a concept of contingent market equilibrium. Recall that an allocation $c \geq 0$ is *feasible* if $\sum_{i=1}^I (c^i - e^i) = 0$. A feasible allocation c is *Pareto efficient* if there is no other feasible allocation \hat{c} such that $\hat{c}^i \succ^i c^i$ for each agent i . Pareto efficiency of an interior allocation is characterized by a nonempty intersection of marginal rates cones:⁹

$$\bigcap_{i=1}^I \nabla U^i[c^i] \neq \emptyset \tag{9}$$

Suppose the sequence of financial markets is replaced with a complete market of contingent claims: There is one contingent claim for each date-event (t, ξ) , which delivers one unit of consumption if the event ξ occurs at date t and nothing in all other date-events. The market price of this claim is the state price $q_t(\xi)$. All affordable consumption plans are collected in the *contingent market budget set*

$$\mathcal{B}^i(q) = \{c^i \in \mathcal{S}_+ \mid q \cdot (c^i - e^i) = 0\}.$$

⁷A proof can be found in the appendix, Lemma 3.

⁸In a two-date economy with nominal assets, Ma (2018) quantifies the dimension of this continuum.

⁹A proof can be found in Rigotti and Shannon (2005), Theorem 3.

Since the market opens only at date 0, all prices are determined simultaneously, price expectations are not necessary, and there can be no dynamic inconsistency. At a contingent market equilibrium, all agents take the state price process $q \in \mathcal{S}$ as given, choose undominated allocations, and the market clears.

Definition 7. A *contingent market equilibrium* for the stochastic economy (\mathcal{I}, Π, u, e) is a tuple (q, c) of state prices and consumption that satisfies

1. for each agent i , $\nexists \hat{c}^i \in \mathcal{B}^i(q)$ such that $\hat{c}^i \succ^i c^i$
2. market clearing $\sum_{i=1}^I c^i - e^i = 0$.

At an interior contingent market equilibrium, the first-order condition

$$q \in \nabla U^i[c^i] \tag{10}$$

is satisfied for each agent i . As this implies (9), every contingent market equilibrium is Pareto efficient.¹⁰ Further, the contingent market budget set \mathcal{B}^i spans the same consumption allocations as the sequence of financial market budget sets (B_0^i, \dots, B_T^i) whenever the market is dynamically complete and agents have correct expectations.¹¹ Recall that $Q[p]$ is a ray in this case. Therefore, the first-order conditions (7) of the financial market at date 0 and (10) of the contingent market agree. As a consequence, every financial market equilibrium has an equivalent contingent market equilibrium, provided that all agents make dynamically consistent decisions.¹² However, a simple example in Section 3 reveals that dynamic consistency is a rather fragile property.

3 Example

Two variations of a particular financial market equilibrium are studied: first, under the maximum-likelihood rule (3), and then, under the full Bayesian rule (2). To isolate the welfare effects of Knightian uncertainty from potential inefficiencies of market incompleteness, the financial market in the example is always dynamically complete. This enables a simple test for Pareto efficiency: Under Assumption 2, preferences are monotone and convex, and in the example, all agents have quasilinear utility functions. Therefore, the second welfare theorem holds in a strong form: *Every Pareto efficient allocation is attained by a contingent market equilibrium and some redistribution of present consumption.* Since the distribution of present consumption does not affect prices, a financial market equilibrium is Pareto efficient if and only if one of its induced state price vectors agrees with the price vector at a contingent market equilibrium.

¹⁰For a general statement that covers boundary equilibria see Rigotti and Shannon (2005), Theorem 1.

¹¹A proof can be found in Kreps (1982), Proposition 2.

¹²In a setting with time-inconsistent preferences, this result is proven in Herings and Rohde (2008).




Figure 1: The date-event tree. Numbers above nodes represent the endowment of Agent 1 (black), numbers below the endowment of Agent 2 (gray). All edges are labeled with the unconditional probability of passing.

3.1 An Example with Well-behaved Equilibria

Consider an economy with three dates and four states of the world $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$. Information unfolds in a binomial tree: The sequence of information partitions is $\mathcal{I}_0 = \{\Omega\}$, $\mathcal{I}_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$, and $\mathcal{I}_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}\}$. This tree is depicted in Figure 1. At date 1, the event $\xi_1 = \{\omega_1, \omega_2\}$ corresponds to the upper branch, and the event $\xi_2 = \{\omega_3, \omega_4\}$ corresponds to the lower branch.

There are $I = 2$ agents with identical utility functions

$$u^i(c^i) = c_0^i + \ln(c_1^i) + \ln(c_2^i) \quad \forall i$$

and identical sets of subjective probabilities

$$\Pi^i = \left\{ \pi^i \in \mathbb{R}_+^{|\Omega|} \mid \pi^i = \left(\frac{\eta^i}{2}, \frac{\eta^i}{2}, \frac{1-\eta^i}{2}, \frac{1-\eta^i}{2} \right), \eta^i \in [0, 1] \right\}.$$

Market outcomes will be studied both in the case of maximum-likelihood updating and in the case of full Bayesian updating. In either case, there is a reason for trade, because the two agents have different endowments:

$$\begin{aligned} e^1 &= \begin{pmatrix} e_0^1(\Omega) & e_1^1(\xi_1) & e_1^1(\xi_2) & e_2^1(\omega_1) & e_2^1(\omega_2) & e_2^1(\omega_3) & e_2^1(\omega_4) \\ 2 & 10/3 & 2/3 & 1/3 & 2/3 & 1/3 & 2/3 \end{pmatrix} \\ e^2 &= \begin{pmatrix} 2 & 2/3 & 10/3 & 2/3 & 1/3 & 2/3 & 1/3 \end{pmatrix} \end{aligned}$$

The financial market consists of $J = 2$ long-lived assets with dividend payments

$$\begin{pmatrix} A^1 \\ A^2 \end{pmatrix} = \begin{pmatrix} A_0(\Omega) & A_1(\xi_1) & A_1(\xi_2) & A_2(\omega_1) & A_2(\omega_2) & A_2(\omega_3) & A_2(\omega_4) \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix}$$

That is, A^1 pays 1 in every upper branch and 0 otherwise, while A^2 pays 1 in every lower branch and 0 otherwise. The easiest way to compute financial market equilibria is to conjecture dynamic completeness and translate contingent market equilibria. The decision criterion of undominatedness is fulfilled by a plan if it maximizes utility under *some* choice of $\eta^i \in [0, 1]$. There is one contingent market equilibrium for each combination of η^1 and η^2 given an arbitrary normalization of prices. If consumption at date 0 is used as the numéraire, the combination $\eta^1 = \eta^2 = 1/2$ results in the state price process

$$q = \begin{pmatrix} q_1(\xi_1) & q_1(\xi_2) & q_2(\omega_1) & q_2(\omega_2) & q_2(\omega_3) & q_2(\omega_4) \\ 1/4 & 1/4 & 1/2 & 1/2 & 1/2 & 1/2 \end{pmatrix}$$

and in consumption plans

$$\begin{pmatrix} C_0^1 \\ C_0^2 \end{pmatrix} = \begin{pmatrix} c_0^i(\Omega) & c_1^i(\xi_1) & c_1^i(\xi_2) & c_2^i(\omega_1) & c_2^i(\omega_2) & c_2^i(\omega_3) & c_2^i(\omega_4) \\ 2 & 2 & 2 & 1/2 & 1/2 & 1/2 & 1/2 \\ 2 & 2 & 2 & 1/2 & 1/2 & 1/2 & 1/2 \end{pmatrix},$$

which are taken as the initial consumption plans at the financial market equilibrium. To translate state prices into asset prices, consider the system of pricing equations

$$\begin{aligned} p_0(\Omega) &= q \cdot A \\ p_1(\xi_1) &= \frac{q_2(\omega_1)}{q_1(\xi_1)} A_2(\omega_1) + \frac{q_2(\omega_2)}{q_1(\xi_1)} A_2(\omega_2) \\ p_1(\xi_2) &= \frac{q_2(\omega_3)}{q_1(\xi_2)} A_2(\omega_3) + \frac{q_2(\omega_4)}{q_1(\xi_2)} A_2(\omega_4). \end{aligned} \tag{11}$$

Let the price expectations P_0 of both agents agree with the solution p to the equation system (11); that is

$$P_0 = \begin{pmatrix} p_0(\Omega) & p_1(\xi_1) & p_1(\xi_2) \\ 5/4 & 5/4 & | & 2 & 2 & | & 2 & 2 \end{pmatrix}.$$

Since asset prices are linearly dependent at each date while dividend payments are linearly independent, the conjecture of dynamic completeness is verified. The consumption plans C_0 are attained by symmetric trading plans

$$\begin{pmatrix} \Psi_0^1 \\ \Psi_0^2 \end{pmatrix} = \begin{pmatrix} \psi_0^i(\Omega) & \psi_1^i(\xi_1) & \psi_1^i(\xi_2) \\ -4/3 & 4/3 & | & 1/6 & -1/6 & | & 1/6 & -1/6 \\ 4/3 & -4/3 & | & -1/6 & 1/6 & | & -1/6 & 1/6 \end{pmatrix},$$




Figure 2: From left to right: Ambiguity partition (solid) before and after observing the event ξ_1 in the original example.

and all markets clear. It is easy to see that C_0 is Pareto efficient: Both agents have identical preferences, and each consumes one-half of the aggregate wealth at each date-event. Will the agents deviate from these plans at date 1? Two cases are considered.

Case 1. Maximum-likelihood updating: Suppose both agents apply the maximum-likelihood rule to update their subjective probabilities. At date 0, there is a maximum of ambiguity about the probabilities of the two events ξ_1 and ξ_2 . However, all this ambiguity vanishes at date 1, once the event is observed. Updating results in a unique probability vector $\Pi_M^i(\xi_1) = \{(1/2, 1/2, 0, 0)\}$ or $\Pi_M^i(\xi_2) = \{(0, 0, 1/2, 1/2)\}$, respectively. For the event ξ_1 , the induced ambiguity partitions are illustrated in Figure 2. The conditional partition $\mathcal{A}^i(\xi_1)$ is finer than the unconditional partition \mathcal{A}^i – *new information reduces uncertainty*.




Figure 3: Section of the budget set before and after the release of information. The marginal rates cone shrinks in a dimension irrelevant for the decision problem. The portfolio decision is thus dynamically consistent.

The consequences of this information release for the decision problem of Agent 1 are illustrated in Figure 3. The left panel displays a section through the initial first-order condition (7) along the $c_2^1(\omega_1)$ – $c_2^1(\omega_2)$ plane. One property should be emphasized: Even

though the marginal rates cone $\nabla U^1[C_0^1]$ is two-dimensional, its section is simply a ray, and it contains the state price vector $(q_2(\omega_1), q_2(\omega_2))$. Although there is Knightian uncertainty at date 0, this plays no role in consumption preferences between state ω_1 and ω_2 . The right panel displays the same section after the information release. There is no change. In fact, the first-order condition (7) at date 0 implies the first-order condition (8) at date 1 because the prior probabilities $\pi^1(\omega_1) = 1/4$ and $\pi^1(\omega_2) = 1/4$ are proportional to the posterior probabilities $\pi^1(\omega_1|\xi_1) = 1/2$ and $\pi^1(\omega_2|\xi_1) = 1/2$. By symmetry of the example, this is true for both agents and both events. Therefore, the original plans (C_0^i, Ψ_0^i) remain optimal. As a consequence, the agents have no reason to revise their price expectations and the realized prices are $P_1 = P_0$, just as expected. The final plans are $C_1 = C_0$ and $\Psi_1 = \Psi_0$, and the tuple (P, C, Ψ) is a financial market equilibrium. All plans are dynamically consistent and *the equilibrium is Pareto efficient*.

Case 2. Full Bayesian updating: Nothing changes if one or both of the agents apply the full Bayesian rule. The resulting posterior probabilities are again $\Pi_{\mathbb{B}}^i(\xi_1) = \{(1/2, 1/2, 0, 0)\}$ and $\Pi_{\mathbb{B}}^i(\xi_2) = \{(0, 0, 1/2, 1/2)\}$. Therefore, full Bayesian updating and maximum-likelihood updating reduce uncertainty in the same way, and in either case, the resulting equilibrium is dynamically consistent and Pareto efficient.

3.2 A Critical Modification of the Example

Consider the following variation of the example. The original set of subjective probabilities is replaced with the following one, which is again shared by both agents:

$$\Pi^i = \left\{ \pi^i \in \mathbb{R}_+^{|\Omega|} \mid \pi^i = \left(\frac{\eta^i}{2}, \frac{1}{4}, \frac{1}{4}, \frac{1-\eta^i}{2} \right), \eta^i \in [0, 1] \right\}.$$

Now, Knightian uncertainty does influence consumption preferences between states ω_1 and ω_2 (resp. ω_3 and ω_4). The natural starting point is again the set of contingent market equilibria. It is parametrized by η^1 and η^2 in the following form:

$$q = \begin{pmatrix} q_1(\xi_1) & q_1(\xi_2) & q_2(\omega_1) & q_2(\omega_2) & q_2(\omega_3) & q_2(\omega_4) \\ \frac{1+\eta^1+\eta^2}{8} & \frac{3-\eta^1-\eta^2}{8} & \frac{\eta^1+\eta^2}{4} & \frac{1}{2} & \frac{1}{2} & \frac{2-\eta^1-\eta^2}{4} \end{pmatrix} \quad (12)$$

Consider again the combination $\eta^1 = \eta^2 = 1/2$, which leads to the same contingent market equilibrium as before. The induced price expectations P_0 and plans (C_0, Ψ_0) are the same as in the original example. The difference manifests only at date 1, when the subjective probabilities are updated.

Case 1. Maximum-likelihood updating: Suppose both agents apply the maximum-likelihood rule; then, updating results in a unique probability vector $\Pi_{\mathbb{M}}^i(\xi_1) = \{(2/3, 1/3, 0, 0)\}$ or $\Pi_{\mathbb{M}}^i(\xi_2) = \{(0, 0, 1/3, 2/3)\}$, respectively. For the event ξ_1 , the induced ambiguity partitions are illustrated in Figure 4. As in the original example, the conditional partition $\mathcal{A}^i(\xi_1)$ is finer than the unconditional partition \mathcal{A}^i , and the information release reduces uncertainty.




Figure 4: From left to right: Ambiguity partition (solid) before and after observing the event ξ_1 in the modified example with maximum-likelihood updating.




Figure 5: Section of the budget set before and after the release of information. The state price vector $(q_2(\omega_1), q_2(\omega_2))$ drops out of the marginal rates cone. The portfolio decision is thus dynamically inconsistent.

However, contrary to the original example, this reduction in uncertainty has an effect, as illustrated in Figure 5 for Agent 1. The left panel shows the first-order condition (7) before the information release: Knightian uncertainty matters and the section of the marginal rates cone $\nabla U^1[C_0^1]$ is two-dimensional. After the information release, the marginal rates cone collapses to a ray, just as in the original example, but now the state price vector is no longer contained. The consequences for the plans of the agent are easily seen: Since there is no uncertainty remaining, the conditional preferences are complete and have an expected utility representation with well-defined indifference surfaces. The indifference surface through the original consumption plan C_0^1 (more precisely, its section on the $c_2^1(\omega_1)$ - $c_2^1(\omega_2)$ plane) is illustrated as the solid curve in the right panel. There is no point of tangency between the curve and the budget set. The optimal allocation involves more consumption in state ω_1 and less consumption in state ω_2 than in the original plan C_0^1 . If both agents adjust their demand, the market cannot clear under the original price expectations P_0 from Section 3.1. Instead, market clearing results in different asset prices

in the date 1 markets:

$$P_1 = \begin{pmatrix} p_0(\Omega) & p_1(\xi_1) & p_1(\xi_2) \\ 5/4 & 5/4 & 8/3 & 4/3 & 4/3 & 8/3 \end{pmatrix}.$$

Both agents react with a new trading plan

$$\begin{aligned} \Psi_1^1 &= \begin{pmatrix} \psi_0^i(\Omega) & \psi_1^i(\xi_1) & \psi_1^i(\xi_2) \\ -4/3 & 4/3 & -1/12 & -5/12 & 5/12 & 1/12 \end{pmatrix} \\ \Psi_1^2 &= \begin{pmatrix} 4/3 & -4/3 & -1/12 & 5/12 & -5/12 & -1/12 \end{pmatrix} \end{aligned}$$

and a new consumption plan

$$\begin{aligned} C_1^1 &= \begin{pmatrix} c_0^i(\Omega) & c_1^i(\xi_1) & c_1^i(\xi_2) & c_2^i(\omega_1) & c_2^i(\omega_2) & c_2^i(\omega_3) & c_2^i(\omega_4) \\ 2 & 1 & 3 & 1/4 & 1/4 & 3/4 & 3/4 \end{pmatrix} \\ C_1^2 &= \begin{pmatrix} 2 & 3 & 1 & 3/4 & 3/4 & 1/4 & 1/4 \end{pmatrix}. \end{aligned}$$

These plans are budget-feasible and clear the market under the revised prices P_1 . As a consequence, the tuple (P, C, Ψ) is a financial market equilibrium. At this equilibrium, both agents make *dynamically inconsistent* decisions. Even though plans, prices, and price expectations are determined at date 0, both agents find it optimal to deviate from their own plans once they reach date 1. Moreover, the equilibrium is qualitatively different: As a result of their original trades at date 0, Agent 1 is wealthier in event ξ_2 while Agent 2 is wealthier in event ξ_1 . To analyze whether this asymmetric distribution of wealth is Pareto efficient, the induced state price process at the equilibrium is computed. This is easily done by solving the equation system (11) for q . Due to dynamic completeness of the financial market, there is a unique solution

$$q = \begin{pmatrix} q_1(\xi_1) & q_1(\xi_2) & q_2(\omega_1) & q_2(\omega_2) & q_2(\omega_3) & q_2(\omega_4) \\ 1/4 & 1/4 & 2/3 & 1/3 & 1/3 & 2/3 \end{pmatrix}.$$

This solution is clearly inconsistent with Equation (12): All contingent market equilibria involve a state price of $1/2$ in states ω_2 and ω_3 . Since there is no equivalent contingent market equilibrium, the dynamically inconsistent financial market equilibrium *is not Pareto efficient*. The source of this inefficiency lies not in market imperfections – these have been ruled out by construction – but in the uncertainty of both agents. At date 0, they are able to choose efficient plans, and they do so. These plans are undominated from the viewpoint of a social planner who aims at efficient risk sharing. However, owing to a lack of dynamic consistency, both agents are tempted to make new plans at date 1. These plans are made under superior information, which eliminates uncertainty and erodes the incentive for risk sharing. From the perspective of the social planner, the resulting new plans are dominated by the original plans. If the agents had adhered to their original plans, no such inefficiency would have occurred.




Figure 6: From left to right: Ambiguity partition (solid) before and after observing the event ξ_1 in the modified example with full Bayesian updating.

Case 2. Full Bayesian updating: It should be noted that information is processed differently under the full Bayesian rule. Contrary to maximum-likelihood updating, the set of posterior probabilities does not become a singleton, but the larger set

$$\Pi_{\text{B}}^i(\xi_1) = \left\{ \pi^i \in \mathbb{R}_+^{|\Omega|} \mid \pi^i = \left(\frac{2\eta^i}{3}, \frac{3-2\eta^i}{3}, 0, 0 \right), \eta^i \in [0, 1] \right\},$$

or its symmetric counterpart if ξ_2 is observed. As a consequence, the dynamic inconsistency illustrated in Figure 5 cannot occur: The state prices $(q_2(\omega_1), q_2(\omega_2))$ coincide with the marginal rates of substitution under the prior probabilities $\pi^1(\omega_1) = 1/4$ and $\pi^1(\omega_2) = 1/4$. While the unique posteriors in $\Pi_{\text{M}}^1(\xi_1)$ are not proportional to these priors, there exists a proportional probability vector in $\Pi_{\text{B}}^1(\xi_1)$. Therefore, the state price vector does not drop out of the marginal rates cone. If both agents apply full Bayesian updating, the resulting financial market equilibrium is dynamically consistent and Pareto efficient.

Unfortunately, the full Bayesian rule leads to a different problem, which is illustrated in Figure 6. The initial ambiguity partition \mathcal{A}^i visualizes the Knightian uncertainty between states ω_1 and ω_4 . There is no such ambiguity regarding the other states: Both agents are able to assign unique probabilities to ω_2 and ω_3 . However, this knowledge is lost once the event ξ_1 is observed. The conditional ambiguity partition $\mathcal{A}^i(\xi_1)$ is not finer than \mathcal{A}^i ; instead, ambiguity has spilled over to the unambiguous event $\{\omega_2\}$. This is inconsistent with the principle that new information should not create new uncertainty – *the agents forget what they knew before*.

To conclude, both updating rules have side effects. Full Bayesian updating may lead to ambiguity spillovers, and agents may fail to recall some of their previous knowledge. Maximum-likelihood updating may lead to dynamic inconsistencies, and agents may abandon Pareto efficient plans. Further, these side effects may occur simultaneously: If one agent in the example applies full Bayesian updating while the other agent applies maximum-likelihood updating, the resulting equilibrium exhibits both ambiguity spillovers and dynamic inconsistencies.

It should be emphasized that these observations are not driven by the selection of a pathological equilibrium. The example is constructed in such a way that full Bayesian updating always results in ambiguity spillovers. This does not depend on a particular equilibrium. The same is true, and this is verified by our general results in Section 4, for maximum-likelihood updating and dynamic inconsistency. Do these updating rules have

desirable properties beside their side effects? And is it possible to design a well-behaved rule that does not come with side effects at all? These questions are addressed in the following section.

4 Results

In a setting with multiple priors, updating is only trivial if there are two dates and all uncertainty is resolved at once. If there are more than two dates, the choice of updating rules has equilibrium effects. This is where decision theory and general equilibrium theory meet: On the one hand, the updating rule of an agent determines his own decisions. On the other hand, it influences the decisions of everyone else in the economy because the choice sets of all agents are connected through the price mechanism. The example from Section 3 gives a hint at this interrelation: Once a single agent changes his plans, other agents respond with a change in their own plans. Thus, dynamically inconsistent behavior spreads through the price mechanism and cannot be predicted by the characteristics of a single agent. Therefore, a characterization of dynamically consistent equilibria cannot be independent of prices. We shall present such a characterization as our first result. For this purpose, we introduce the *supporting probabilities* of a state price process q at the allocation c^i :

$$\bar{\pi}^i[q, c^i] = \{ \pi^i \in \Pi^i \mid q = \alpha(\pi^i \otimes Du^i[c^i]) \text{ for some } \alpha > 0 \} \quad (13)$$

Note the following properties: Since $\Pi^i \subseteq \Delta^{|\Omega|}$, the set $\bar{\pi}^i[q, c^i]$ is either a singleton or empty. It is nonempty for some $q \in Q[p]$ if the first-order condition (7) is satisfied. It is always empty if the condition is not satisfied. If there is no Knightian uncertainty or if the financial market is dynamically complete, each agent has a unique supporting probability vector at any equilibrium. If Knightian uncertainty meets market incompleteness, an equilibrium can be supported by multiple subjective probability vectors. The following theorem presents a necessary and sufficient condition for dynamically consistent equilibria: One supporting probability vector must be contained in the convex hull of all updates, date by date.

Theorem 1. *Let (\mathcal{I}, Π, u, e) be a stochastic economy with asset structure A . An interior financial market equilibrium (P, C, Ψ) is dynamically consistent if and only if there is some $q \in Q[P_0]$ such that for each agent i ,*

$$\bar{\pi}^i[q, C_0^i] \in \text{conv} \left((\Pi^i(\xi))_{\xi \in \mathcal{I}_t} \right) \quad \forall t \geq 0.$$

Theorem 1 is a particularly strong result if the financial market is dynamically complete. In this case, $Q[P_0]$ is a singleton and the test for dynamic consistency boils down to checking the unique supporting probability vector. In addition, the theorem offers insights on the interplay of market incompleteness and dynamic consistency well beyond the dynamically complete case. If the market becomes more and more incomplete, fewer and fewer economies permit dynamically inconsistent decisions. The reason is that the

dimension of $Q[P_0]$ grows with the degree of market incompleteness. In the extreme case of an asset structure without assets, $\dim(Q[P_0]) = |\Omega|$ and the theorem implies that all equilibria are dynamically consistent – if no trade is the only available trading plan, no agent can ever switch to a different plan.

The implications of the theorem are more interesting if the asset structure is not degenerate: Even in this general case, there exists a nontrivial subset of economies in which all equilibria are dynamically consistent; for instance, the economy from the original example in Section 3. It is irrelevant whether agents apply maximum-likelihood updating, full Bayesian updating, or some other updating rule – any rule Π^i consistent with Bayes' law satisfies $\eta^i \Pi^i(\xi_1) + (1 - \eta^i) \Pi^i(\xi_2) = \eta^i (1/2, 1/2, 0, 0) + (1 - \eta^i) (0, 0, 1/2, 1/2) = \Pi^i$. A simple application of Theorem 1 shows that all equilibria are dynamically consistent.

In the modified version of the example, some equilibria are dynamically inconsistent if at least one agent applies the maximum-likelihood rule. This observation can be explained by means of the following corollary; in fact, the corollary strengthens the finding from the example.

Corollary 1. *Let (\mathcal{I}, Π, u, e) be a stochastic economy with asset structure A . All interior financial market equilibria are dynamically inconsistent if for some agent i and date $t > 0$*

$$\Pi^i \cap \text{conv} \left((\Pi^i(\xi))_{\xi \in \mathcal{I}_t} \right) = \emptyset.$$

Proof. By construction $\bar{\pi}^i[q, C_0^i] \in \Pi^i$. If Π^i does not intersect the convex hull, the condition from Theorem 1 is violated. \square

Since the convex hull $\eta^i \Pi_M^i(\xi_1) + (1 - \eta^i) \Pi_M^i(\xi_2) = \eta^i (2/3, 1/3, 0, 0) + (1 - \eta^i) (0, 0, 1/3, 2/3)$ does not intersect the set of prior probabilities Π^i , it follows from Corollary 1 that all equilibria are dynamically inconsistent.

As far as positive theory is concerned, these findings can be summarized as follows: *Depending on the updating rule, dynamic consistency is prevalent in some economies while dynamic inconsistency is prevalent in others.* As regards normative theory, it is not clear per se whether dynamically inconsistent decisions should be viewed as undesirable. After all, agents are simply reacting to new information. To make a normative judgment, one must define criteria of social optimality and individual rationality, and check whether these are fulfilled, possibly under a suitable restriction of updating rules.

4.1 Full Bayesian Updating

The most common criterion of social optimality is Pareto efficiency. In a dynamically complete market this criterion is attainable because arbitrary transfers of future consumption are feasible. In the example from Section 3, Pareto efficient transfers occur in equilibrium if all agents apply the full Bayesian rule (2): Initially, all agents make Pareto efficient plans, just as in the contingent market equilibrium of Rigotti and Shannon (2005). Later, and this is where the updating rule matters, no adjustments are made and the plans are dynamically consistent. The first step toward Pareto efficiency is therefore to establish dynamic consistency as a general property:

Corollary 2. *Let (\mathcal{I}, Π, u, e) be a stochastic economy with asset structure A . All interior financial market equilibria are dynamically consistent if each agent i applies the full Bayesian rule:*

$$\Pi^i(\xi) = \Pi_{\mathbf{B}}^i(\xi) \quad \forall \xi \in \mathcal{I}_t, t > 0.$$

Proof. By Equation (2),

$$\pi^i \in \Pi^i \quad \Longrightarrow \quad \frac{\text{pr}_{\xi}(\pi^i)}{\pi^i(\xi)} \in \Pi_{\mathbf{B}}^i(\xi) \quad \forall \xi \in \mathcal{I}_t, t \geq 0. \quad (14)$$

Define a vector of weights $\alpha \in \Delta_+^{|\mathcal{I}_t|}$ componentwise as $\alpha_{\xi} = 1/\pi^i(\xi)$. By construction of the projection pr_{ξ} , the prior probability vector is $\pi^i = \sum_{\xi \in \mathcal{I}_t} \text{pr}_{\xi}(\pi^i)$. It follows from (14) that

$$\pi^i \in \Pi^i \quad \Longrightarrow \quad \pi^i \in \sum_{\xi \in \mathcal{I}_t} \alpha_{\xi} \Pi_{\mathbf{B}}^i(\xi) \quad \forall t \geq 0,$$

and thus

$$\Pi^i \subseteq \text{conv} \left(\left(\Pi_{\mathbf{B}}^i(\xi) \right)_{\xi \in \mathcal{I}_t} \right) \quad \forall t \geq 0.$$

Dynamic consistency follows from Theorem 1 because $\bar{\pi}^i[q, C_0^i] \in \Pi^i$. \square

Corollary 2 reveals a remarkable difference between the Knightian decision model and maxmin expected utility when it comes to sequential decision making: It is well-known that full Bayesian updating is not sufficient for dynamic consistency under maxmin expected utility. The only exceptions are decision problems that satisfy the rectangularity condition of Epstein and Schneider (2003). In all other cases, dynamic consistency is an unrealistic requirement, and not simply a deficiency of full Bayesian updating: As Hanany and Kilbanoff (2007) show, there exists no updating rule consistent with Definition 1 that guarantees dynamically consistent decisions in the maxmin expected utility setting. By contrast, in the present setting such an updating rule exists, and it is identified in the corollary. Further restrictions, such as rectangularity, are not necessary.

Since all agents in the economy are required to apply the full Bayesian rule, Corollary 2 is a rather demanding statement. However, the corollary cannot be reduced to a statement of individual behavior. An agent who applies the full Bayesian rule would not react to new information with a change of plans *ceteris paribus*. However, if prices and price expectations adjust to a change in demand of agents with other updating rules, even the full Bayesian updater will make dynamically inconsistent decisions. If this is ruled out, Pareto efficiency follows in a second step:

Corollary 3. *Let (\mathcal{I}, Π, u, e) be a stochastic economy with asset structure A . All interior financial market equilibria with a dynamically complete market are Pareto efficient if each agent i applies the full Bayesian rule:*

$$\Pi^i(\xi) = \Pi_{\mathbf{B}}^i(\xi) \quad \forall \xi \in \mathcal{I}_t, t > 0.$$

Proof. At any equilibrium, the date 0 first-order condition (7) is fulfilled for each consumer i . By Corollary 2, full Bayesian updating implies dynamic consistency, and by Lemma 3, $Q[P_0]$ is one-dimensional. As a consequence, (7) implies the efficiency condition (9). Thus, $C_0 = \dots = C_T$ is a Pareto efficient allocation. \square

Corollary 3 shows that dynamic consistency is a sufficient condition for Pareto efficiency, provided the market is not too incomplete. In the example from Section 3, it seems as if this condition was also necessary. Can the corollary be strengthened into a characterization? The answer is in the negative, because preferences are incomplete. It is true that dynamic inconsistency implies that agents move away from an efficient allocation. However, the Pareto frontier is a large set. It may well be that the new allocation is still contained in the Pareto frontier, simply because old and new consumption plans are not comparable for some agents.

Even though Corollary 3 may appear restrictive as it only holds for dynamically complete financial markets, it is the most general welfare result one can hope to obtain. Without dynamic completeness, Pareto efficiency is out of reach because the necessary transfers are typically not feasible in the financial market. As Ma (2015) shows, the weaker standard of constrained efficiency is still attainable if the financial market opens only once. However, if the market opens sequentially, like in the present model, results of Geanakoplos and Polemarchakis (1986) and Citanna, Kajji, and Villanacci (1998) indicate that not even constrained efficiency can be expected. In the single-prior case, constrained inefficiency is a generic phenomenon, and it remains robust when agents have multiple priors. Therefore, the corollary cannot be extended outside the setting with dynamically complete markets. Within this setting, however, it offers a strong justification for full Bayesian updating. The rule leads to a Pareto efficient market outcome and thus fulfills a normative criterion of social optimality.

4.2 Maximum-likelihood Updating

The criterion of individual rationality we are most concerned with is that agents do not forget what they knew. In the single-prior setting, this simply means once the agent assigns a probability of zero to an event, this probability can never again become positive. This condition is fulfilled as long as the information partitions become progressively finer and the updates satisfy Bayes' law. In the multi-prior setting, a stronger condition is necessary. A minimal requirement is that once the agent assigns a unique probability to an event, this probability can never again become ambiguous. This is ensured if the ambiguity partitions become progressively finer. However, this condition is not yet strong enough to guarantee that updates are consistent with the agent's previous knowledge.

The issue shall be illustrated in a small example. The focus of the example is on the decision problem of a single agent. There is Knightian uncertainty, but this uncertainty is not payoff-relevant. There is an information release, but this information release is also not payoff-relevant. Nevertheless, the following phenomenon occurs under some updating rules: Initially, the agent assigns unique probabilities to all payoff-relevant events. However, just as in the example from Section 3, new information causes ambiguity spillovers,




Figure 7: Problematic full Bayesian updating: The marginal distribution of the events $(\zeta_1, \zeta_2, \zeta_3)$ changes in spite of no news about their probabilities.

and the agent becomes uncertain about the probabilities he knew before. If we set unambiguous probabilities equal to objective probabilities, as in Bewley (1986, 2002), this means that the agent starts with a correct model of the world, but discards this model after receiving irrelevant information.

Example. Consider a state space with six states of the world, $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$, which occur with equal probabilities of $1/6$. The agent has incomplete knowledge of these objective probabilities and perceives Knightian uncertainty: He can identify three unambiguous events,

$$\zeta_1 = \{\omega_1, \omega_2\}, \quad \zeta_2 = \{\omega_3, \omega_4\}, \quad \text{and} \quad \zeta_3 = \{\omega_5, \omega_6\}.$$

Each of these occurs with a known probability of $\pi^i(\zeta_1) = \pi^i(\zeta_2) = \pi^i(\zeta_3) = 1/3$. However, there is complete ambiguity about the probabilities of the two states within each of the three events. This ambiguity partition is illustrated in the left panel of Figure 7. The agent has to choose between a risk-free asset $A^1 = \vec{1}$, whose payoffs are 1 in each state, and a risky asset $A^2 = (1/2, 1/2, 1, 1, 2, 2)$. Note that only the marginal distribution of $(\zeta_1, \zeta_2, \zeta_3)$ matters. Since the agent knows this marginal distribution, ambiguity plays no role in the decision problem and, as long as preferences satisfy Assumption 2, there is only one undominated choice. However, suppose the agent receives the information that the true state of the world is contained in $\xi = \{\omega_2, \omega_4, \omega_6\}$ before he makes his decision. As a consequence of this release, all odd states can be ruled out. This release is irrelevant since it contains no news about the marginal distribution. The only consistent one is still:

$$\pi^i(\zeta_1|\xi) = \pi^i(\zeta_2|\xi) = \pi^i(\zeta_3|\xi) = 1/3.$$

However, if the agent applies the full Bayesian rule, he does not reach this conclusion. Updating Π^i prior by prior results in

$$\Pi_{\mathbb{B}}^i(\xi) = \left\{ \pi^i \in \mathbb{R}_+^{|\Omega|} \mid \pi^i(\omega_2) + \pi^i(\omega_4) + \pi^i(\omega_6) = 1 \right\},$$

and all of a sudden all choices are undominated. The ambiguity spillover is illustrated in the right panel of Figure 7. This spillover may lead to completely different decisions even though the observed information is not payoff-relevant.

Such updating is problematic since it adds something arbitrary to the set of posteriors. Consider, for example, the probability vector $\pi^i = (1, 0, 0, 0, 0, 0)$, which puts all the mass

on state ω_1 . This vector is a valid posterior under full Bayesian updating, but the marginal distribution has changed completely. Even though the information set ξ does not rule out the events ζ_2 and ζ_3 , the agent finds it possible that both occur with a probability of zero. This is inconsistent with his priors, which assign an unambiguous probability of $1/3$ to both events. The issue with the updating rule in the example can be described as follows. Full Bayesian updating allows ambiguous probability mass to flow from one event to another. Although the new information resolves all uncertainty, the probability mass is no longer distributed in equal proportions over the three events.

How do we avoid this kind of behavior? It is not sufficient to impose a restriction on the progressive fineness of ambiguity partitions. Consider, for example, an updating rule that results in $\Pi^i(\xi) = \{(1, 0, 0, 0, 0)\}$. Such an update is consistent with Bayes' law and leads to a finer ambiguity partition. Nevertheless, it is at odds with the previous knowledge of the agent – he forgets the objective probabilities of the payoff-relevant events. To prevent such misguided information processing, the updating rule must preserve the marginal distribution of all events that are still considered possible after an information release. In other words, their likelihood ratios must not change between priors $\tilde{\pi}^i$ and posteriors π^i :

$$\frac{\pi^i(\xi \cap \zeta)}{\pi^i(\xi \cap \zeta')} = \frac{\tilde{\pi}^i(\zeta)}{\tilde{\pi}^i(\zeta')} \quad \forall \zeta, \zeta' \in \mathcal{A}^i, \xi \cap \zeta \neq \emptyset, \xi \cap \zeta' \neq \emptyset$$

This condition is the necessary refinement of full Bayesian updating that eliminates the undesirable behavior in the above example. It forces ambiguous probability mass to stay within a cell of \mathcal{A}^i as long as some states are still considered possible. In the above example, this refinement works well because only ambiguous probability mass has to be redistributed. However, if there is unambiguous probability mass on some states in a cell of \mathcal{A}^i , this mass cannot remain within the cell when those states are discarded; otherwise, Bayes' law would be violated. For any event ξ , the unambiguous probability mass is $\min_{\hat{\pi}^i \in \Pi^i} \hat{\pi}^i(\xi)$, the least probability that all priors assign. To maintain Bayes' law, unambiguous mass on discarded states has to be subtracted when calculating the likelihood ratios:

$$\frac{\pi^i(\xi \cap \zeta)}{\pi^i(\xi \cap \zeta')} = \frac{\tilde{\pi}^i(\zeta) - \min_{\hat{\pi}^i \in \Pi^i} \hat{\pi}^i(\zeta \setminus \xi)}{\tilde{\pi}^i(\zeta') - \min_{\hat{\pi}^i \in \Pi^i} \hat{\pi}^i(\zeta' \setminus \xi)} \quad \forall \zeta, \zeta' \in \mathcal{A}^i, \xi \cap \zeta \neq \emptyset, \xi \cap \zeta' \neq \emptyset.$$

This is the general condition we impose to ensure that agents do not lose their previous knowledge. Since $\tilde{\pi}^i(\zeta) = \min_{\hat{\pi}^i \in \Pi^i} \hat{\pi}^i(\zeta \setminus \xi) + \max_{\hat{\pi}^i \in \Pi^i} \hat{\pi}^i(\xi \cap \zeta)$ for any $\zeta \in \mathcal{A}^i$, it can be stated more concisely in the form of the following definition.

Definition 8. An updating rule satisfies *ambiguity containment* if for any $\xi \subseteq \Omega$ the following condition is true:

$$\pi^i \in \Pi^i(\xi) \implies \frac{\pi^i(\xi \cap \zeta)}{\pi^i(\xi \cap \zeta')} = \frac{\max_{\hat{\pi}^i \in \Pi^i} \hat{\pi}^i(\xi \cap \zeta)}{\max_{\hat{\pi}^i \in \Pi^i} \hat{\pi}^i(\xi \cap \zeta')} \quad \forall \zeta, \zeta' \in \mathcal{A}^i, \xi \cap \zeta' \neq \emptyset.$$

Even though the space of potential updating rules is very large, it turns out that the conditions from Definitions 1 and 8 restrict this space to one single point. Bayes' law and ambiguity containment are jointly satisfied if and only if the maximum-likelihood rule (3) is used:

Proposition 1. *The unique updating rule that satisfies Bayes' law and ambiguity containment is the maximum-likelihood rule:*

$$\Pi^i(\xi) = \Pi_{\mathbb{M}}^i(\xi) \quad \forall \xi \subseteq \Omega.$$

Proposition 1, whose proof is presented in the appendix, offers a strong justification for maximum-likelihood updating. It shows that under any other rule, agents may lose confidence in their own model of the world on receiving irrelevant information. The maximum-likelihood rule ensures that new information does not create new uncertainty and thus fulfills a normative criterion of individual rationality.

Contrary to full Bayesian updaters, agents who apply the maximum-likelihood rule may revise their original plans. It is worth noting that this creates a propensity to trade when the financial market reopens, even under a high degree of uncertainty. This stands in contrast to the two-date model of Rigotti and Shannon (2005), in which trade breaks down if uncertainty grows large. In extreme cases with prevalent uncertainty, no trade may be the only individually rational contingent market equilibrium. However, in a financial market that opens sequentially, the presence of a single maximum-likelihood updater may be enough to eliminate no-trade equilibria. Under a weak condition on his subjective probabilities, there are always economies in which all equilibria involve trade, *no matter what the characteristics of the other agents are.*

Corollary 4. *Let $|\Omega| \geq 4$, $J \geq 1$. If some agent i applies the maximum-likelihood rule,*

$$\Pi^i(\xi) = \Pi_{\mathbb{M}}^i(\xi) \quad \forall \xi \subseteq \Omega,$$

and if at least two states are unambiguous while others are ambiguous, then there exists an information structure \mathcal{I} such that any stochastic economy (\mathcal{I}, Π, u, e) has the following property: All financial market equilibria involve trade.

Proof. Without loss of generality, let ω_1 and ω_2 be the unambiguous states. Then there exists an information partition $\mathcal{I}_1 = \{\xi_1, \xi_2\}$ with $\{\omega_1\} \subset \xi_1$ and $\{\omega_2\} \subset \xi_2$. Under Assumption 1, $\pi^i(\xi_1) > 0$ and $\pi^i(\xi_2) > 0$ for any $\pi^i \in \Pi^i$. The presence of ambiguity in Π^i implies that

$$\pi^i(\xi_1) + \hat{\pi}^i(\xi_2) \neq 1 \quad \forall \pi^i, \hat{\pi}^i \in \Pi^i, \pi^i \neq \hat{\pi}^i. \quad (15)$$

The corollary can now be proven by contradiction: Let $\pi^i(\cdot|\xi_1) \in \Pi_{\mathbb{M}}^i(\xi_1)$ and $\hat{\pi}^i(\cdot|\xi_2) \in \Pi_{\mathbb{M}}^i(\xi_2)$; then, $\pi^i \neq \hat{\pi}^i$ by definition of the rule (3). Suppose that some convex combination satisfies $\alpha\pi^i(\cdot|\xi_1) + (1-\alpha)\hat{\pi}^i(\cdot|\xi_2) \in \Pi^i$. Since ω_1 and ω_2 are unambiguous, this implies $\alpha\pi^i(\omega_1|\xi_1) = \pi^i(\omega_1)$ and $(1-\alpha)\hat{\pi}^i(\omega_2|\xi_2) = \hat{\pi}^i(\omega_2)$. Applying (1) to these two identities results in $\alpha = \pi^i(\xi_1)$ and $1-\alpha = \hat{\pi}^i(\xi_2)$, but this violates (15) and is therefore impossible. Thus, the maximum-likelihood rule implies $\Pi^i \cap \text{conv}(\Pi_{\mathbb{M}}^i(\xi_1), \Pi_{\mathbb{M}}^i(\xi_2)) = \emptyset$ and Corollary 1

can be invoked: Even if the initial plan involves no trade, dynamic inconsistency implies that agents start trading at date 1. Since Assumption 3 rules out endowments at the boundary, no trade implies interiority and boundary equilibria need not be considered. Only if $J = 0$ do budget sets contain no interior points, but this case is excluded. \square

Corollary 4 is empirically well-founded. It holds as long as *some* agents use maximum-likelihood updating, and this is supported by experimental results. In a dynamic version of the Ellsberg urn experiment, Cohen, Gilboa, Jaffray, and Schmeidler (2000) classify about 29% of subjects as maximum-likelihood updaters. In a simulated asset market experiment with an information structure identical to the one from the example in Section 3.2, Ngangoué (2018) identifies the behavior of 22% of subjects as consistent with maximum-likelihood updating, and these subjects always trade. The results become stronger if one tests directly for dynamic consistency: Again in a dynamic Ellsberg urn experiment, Dominiak, Duersch, and Lefort (2012) find that about 68% of subjects make dynamically inconsistent decisions.

4.3 A Difficulty in the Design of Updating Rules

Now that normative criteria of social optimality and individual rationality are formulated, it remains to be analyzed when and whether these criteria are fulfilled in equilibrium. Taken separately, neither of the two criteria is too demanding. There are choices of updating rules that ensure either: Full Bayesian updating guarantees dynamic consistency, which translates into Pareto efficiency in dynamically complete markets (Corollary 3). Maximum-likelihood updating guarantees ambiguity containment, and that agents do not forget (Proposition 1). In light of these findings, a natural desideratum of normative theory is to design an updating rule that combines both criteria. Unfortunately, this is impossible:

Corollary 5. *There exists no updating rule that satisfies Bayes' law and guarantees both dynamically consistent plans and ambiguity containment.*

Proof. By Proposition 3, the only rule that jointly satisfies Bayes' law and ambiguity containment is the maximum-likelihood rule. However, in Section 3, it is demonstrated that this rule may lead to dynamically inconsistent plans. \square

Corollary 5 highlights a conflict between normative criteria under Knightian uncertainty. In simple two-date models, this conflict does not surface because all uncertainty is eliminated at once. In this case, all updating rules agree and assign a probability of one to the realized state. For the same reason, conflicts of this kind do not occur under subjective expected utility: If there is a single prior, all updating rules agree and result in simple Bayesian updating. In this case, plans are always dynamically consistent and new information always reduces uncertainty. Difficulties only appear when the updating rules are in discord.

Even though these difficulties cannot be resolved in general, the insights from Section 3 suggest that some economies are unproblematic in that no conflict between the

two normative criteria occurs. This is the case whenever the full Bayesian rule and the maximum-likelihood rule lead to equivalent updates. Our next result is a simple characterization of unproblematic economies. It depends only on the relative fineness of information partitions and ambiguity partitions. If an observed event ξ shares states with a cell ζ of the ambiguity partition, the rules agree if $\xi \subseteq \zeta$; that is, if the event is fully contained in a cell of the ambiguity partition. The same holds true if $\zeta \subseteq \xi$. The following theorem shows that these easy-to-verify criteria characterize economies in which dynamic consistency and ambiguity containment are compatible.

Proposition 2. *The following statements are equivalent:*

(i) For each $t > 0$,

$$\xi \cap \zeta \in \{\xi, \zeta, \emptyset\} \quad \forall (\xi, \zeta) \in \mathcal{I}_t \times \mathcal{A}^i.$$

(ii) For each $t > 0$,

$$\Pi_{\mathbb{B}}^i(\xi) = \Pi_{\mathbb{M}}^i(\xi) \quad \forall \xi \in \mathcal{I}_t.$$

The aesthetics of Proposition 2 is in its generality. It is based on a criterion for economies and therefore independent of equilibrium variables. If the criterion is satisfied, then *all* equilibria are dynamically consistent and no ambiguity spillovers occur, regardless of how incomplete the asset structure A is and *regardless of which updating rules the agents apply*.

Finally, we relate our criterion to the rectangularity condition of Epstein and Schneider (2003), which characterizes dynamically consistent decisions under full Bayesian updating in the related setting of maxmin expected utility. In the notation of the present paper, the condition can be stated as follows. Summations are understood as Minkowski sums.

Definition 9. A tuple (\mathcal{I}, Π^i) of information structure and subjective probabilities satisfies *rectangularity* if for each $t > 0$ and each $\xi \in \mathcal{I}_{t-1}$,

$$\Pi_{\mathbb{B}}^i(\xi) = \bigcup_{\pi^i \in \Pi^i} \sum_{\xi' \in \mathcal{I}_t} \pi^i(\xi' | \xi) \Pi_{\mathbb{B}}^i(\xi').$$

Since Definition 9 is based on the full Bayesian rule, $\Pi_{\mathbb{B}}^i$, it may create the impression that full Bayesian updating and rectangularity are intrinsically related. However, the condition is in fact so strong that differences between updating rules disappear.

Proposition 3. *If (\mathcal{I}, Π^i) satisfies rectangularity, then for each $t > 0$,*

$$\Pi_{\mathbb{B}}^i(\xi) = \Pi_{\mathbb{M}}^i(\xi) \quad \forall \xi \in \mathcal{I}_t.$$

Proposition 3 shows that rectangularity is at least as strong as the criterion from Proposition 2. Since this criterion is less obscure, there is no value added in adopting the rectangularity condition within the framework of Knightian decision theory. In fact, our criterion is strictly weaker than rectangularity, which is easily seen in the final example.

Example. Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ be the state space and let $\mathcal{I}_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}\}$ be the date 1 information partition. Consider an agent with subjective probabilities $\Pi^i = \{\pi^i \in \Delta^3 \mid \pi^i(\omega_1) \leq 1/2\}$. As the induced ambiguity partition is $\mathcal{A}^i = \{\Omega\}$, the condition from Proposition 2 is satisfied whereas rectangularity is violated.

5 Conclusion

In the traditional theory of general equilibrium, financial market models have two important properties: First, all consumption and trading plans are dynamically consistent, and second, better information reduces uncertainty. If Knightian uncertainty is added to the model, these two properties are at odds. The reason is that both properties depend on the updating rule for subjective probabilities. Even though we have identified rules that guarantee either property, there exists no rule that guarantees both. This impossibility result poses a challenge for a general theory of dynamic markets under Knightian uncertainty. One of the two properties must be given up, but neither is a natural candidate.

Allowing uncertainty to spread is not only at odds with traditional theory but incompatible with the ideas of Knight (1921). Even though he employs his own concept of uncertainty, he conforms with the traditional view that uncertainty stems from imperfect knowledge. In line with this view, he recognizes that uncertainty is reduced through new information. This is expressed most concisely in his discussion of social aspects of uncertainty: “*The amount of uncertainty may, however, be reduced in several ways, as we have seen. In the first place, we can increase our knowledge of the future through scientific research and the accumulation and study of the necessary data*” (p. 347). The argument is normative in nature: Agents should use information to improve their knowledge, and should not forget knowledge previously held.

Dynamic consistency, on the other hand, has its own normative implication: It guarantees a Pareto efficient market outcome, provided the market is dynamically complete. If dynamic consistency is given up, Pareto efficiency is lost in some economies. It must be noted, though, that there is a set of economies in which both properties are compatible because different updating rules lead to identical updates. Our characterization shows that this set is sizable. In particular, it contains the well-known special cases of two-date economies and single-prior economies. Outside this set, however, different updating rules have different implications. In this case, the theory’s predictions depend on how agents update subjective probabilities, which is ultimately an empirical question.

Appendix

The appendix contains the proofs of Theorem 1 and all propositions. These proofs make use of the following two lemmata:

Lemma 1. *Let $\Pi^i : 2^\Omega \rightrightarrows \Delta^{|\Omega|}$ be an updating rule that satisfies Bayes’ law. Then, for any partition \mathcal{I}_t of Ω ,*

$$\text{pr}_\xi(q) \in \text{cone}(\Pi^i(\xi)) \quad \forall \xi \in \mathcal{I}_t \quad \iff \quad q \in \text{cone} \left((\Pi^i(\xi))_{\xi \in \mathcal{I}_t} \right).$$

Proof. Necessity: If q belongs to the conical hull, there is some vector of weights $\alpha \in \mathbb{R}_+^{|\mathcal{I}_t|}$ such that $q \in \sum_{\xi \in \mathcal{I}_t} \alpha_\xi \Pi^i(\xi)$. For any $\xi \in \mathcal{I}_t$, the projection pr_ξ can be applied to both sides, which yields

$$\text{pr}_\xi(q) \in \alpha_\xi \text{pr}_\xi(\Pi^i(\xi)) \quad \forall \xi \in \mathcal{I}_t,$$

and since $\text{pr}_\xi(\Pi^i(\xi)) = \Pi^i(\xi)$ by Bayes' law, this implies the left-hand side of the equivalence.

Sufficiency: By construction of the projection pr_ξ , any vector q that satisfies the left-hand side of the equivalence can be recovered as $q = \sum_{\xi \in \mathcal{I}_t} \text{pr}_\xi(q)$. Therefore, it satisfies

$$q \in \sum_{\xi \in \mathcal{I}_t} \alpha_\xi \text{cone}(\Pi^i(\xi))$$

with $\alpha = \vec{1}$. This means that q belongs to the conical hull. \square

Lemma 2. *For any event $\xi \subseteq \Omega$, there exists a probability vector $\pi^{i*} \in \Pi^i$ that satisfies*

$$\pi^{i*} \in \arg \max_{\pi^i \in \Pi^i} \pi^i(\xi \cap \zeta) \quad \forall \zeta \in \mathcal{A}^i \text{ with } \xi \cap \zeta \text{ nonempty},$$

and this condition is equivalent to

$$\pi^{i*} \in \arg \max_{\pi^i \in \Pi^i} \pi^i(\xi).$$

Proof. As regard existence, note that Π^i is compact under Assumption 1 as a closed subset of the unit simplex. Thus, the final problem $\max_{\pi^i \in \Pi^i} \pi^i(\xi)$ has a solution. As regards equivalence, the implication from the first condition to the second is trivial. The reverse implication can be proven by contradiction: Any maximizer π^{i*} from the second condition also solves the maximization problems in the first condition. Suppose not; then, there would be some $\zeta \in \mathcal{A}^i$ and some $\pi^i \in \Pi^i$ such that $\pi^{i*}(\xi \cap \zeta) < \pi^i(\xi \cap \zeta)$. By construction of \mathcal{A}^i , π^i can be chosen to satisfy $\pi^i(\zeta') = \pi^{i*}(\zeta')$ for all $\zeta' \neq \zeta$ because $\Pi^i = \sum_{\zeta \in \mathcal{A}^i} \text{pr}_\zeta(\Pi^i)$. Summing over all $\zeta' \in \mathcal{A}^i$ with $\xi \cap \zeta'$ nonempty leads to

$$\pi^{i*}(\xi) = \sum_{\zeta'} \pi^{i*}(\xi \cap \zeta') < \sum_{\zeta' \neq \zeta} \pi^{i*}(\xi \cap \zeta') + \pi^i(\xi \cap \zeta) = \pi^i(\xi),$$

but if $\pi^{i*}(\xi) < \pi^i(\xi)$, then π^{i*} could not have been a maximizer in the first place. \square

Let the operator \oslash denote the inverse of the componentwise product \otimes , the componentwise division: $(Y \otimes x) \oslash x = Y$. It satisfies $\text{pr}_\xi(y) \oslash x = \text{pr}_\xi(y \oslash x)$. The componentwise product on \mathcal{S} can also be defined between a set $Y \subseteq \mathcal{S}$ and a vector $x \in \mathcal{S}$ as $Y \otimes x = \{z \in \mathcal{S} \mid z = x \otimes y \text{ for some } y \in Y\}$.

Proof of Theorem 1. Dynamic consistency means that for each agent i , the original consumption plan satisfies all first-order conditions; i.e.,

$$Q[P_0](\xi) \cap \nabla U^i[C_0^i](\xi) \neq \emptyset \quad \forall \xi \in \mathcal{I}_t, t \geq 0. \quad (16)$$

By definition, $Q[P_0](\xi) = \text{pr}_\xi(Q[P_0])$, and by Equation (6), $\nabla U^i[C_0^i](\xi) = \text{cone}(\Pi^i(\xi)) \otimes Du^i[C_0^i]$. Using these facts, (16) is equivalently expressed as

$$\exists q \in Q[P_0] \text{ such that } \text{pr}_\xi(q) \in \text{cone}(\Pi^i(\xi)) \otimes Du^i[C_0^i] \quad \forall \xi \in \mathcal{I}_t, t \geq 0. \quad (17)$$

Applying \circledast to both sides of (17) yields the equivalent expression

$$\exists q \in Q[P_0] \text{ such that } \text{pr}_\xi(q \circledast Du^i[C_0^i]) \in \text{cone}(\Pi^i(\xi)) \quad \forall \xi \in \mathcal{I}_t, t \geq 0. \quad (18)$$

Moreover, if \circledast is applied to Equation (13), it becomes clear that $q \circledast Du^i[C_0^i] = \alpha \bar{\pi}^i[q, C_0^i]$ for some $\alpha > 0$. Combining this with (18) results in the equivalent expression

$$\exists q \in Q[P_0] \text{ such that } \text{pr}_\xi(\bar{\pi}^i[q, C_0^i]) \in \text{cone}(\Pi^i(\xi)) \quad \forall \xi \in \mathcal{I}_t, t \geq 0, \quad (19)$$

and by Lemma 1, (19) is equivalent to

$$\exists q \in Q[P_0] \text{ such that } \bar{\pi}^i[q, C_0^i] \in \text{cone} \left((\Pi^i(\xi))_{\xi \in \mathcal{I}_t} \right) \quad \forall t \geq 0. \quad (20)$$

Finally, note that $\bar{\pi}^i[q, C_0^i] \in \Delta^{|\Omega|}$ and $\Pi^i(\xi) \subseteq \Delta^{|\Omega|}$. Therefore, attention can be restricted to the unit simplex and the conical hull in (20) can be replaced with the convex hull, which results in the condition of the theorem. \square

Proof of Proposition 1. Necessity of Π_M^i : By Lemma 2, the condition of ambiguity containment can be written as

$$\frac{\pi^i(\xi \cap \zeta)}{\pi^i(\xi \cap \zeta')} = \frac{\pi^{i*}(\xi \cap \zeta)}{\pi^{i*}(\xi \cap \zeta')} \quad \forall \zeta, \zeta' \in \mathcal{A}^i \text{ with } \xi \cap \zeta' \neq \emptyset$$

which defines a system of equations. Although the system is overdetermined, there exists some scalar $\mu \neq 0$ such that

$$\pi^i(\xi \cap \zeta) = \mu \pi^{i*}(\xi \cap \zeta) \quad \forall \zeta \in \mathcal{A}^i$$

constitutes a solution. By definition, $\pi^i(\xi) = 1$ is satisfied for $\pi^i \in \Pi^i(\xi)$ under any updating rule $\Pi^i(\xi)$. Therefore,

$$\begin{aligned} 1 = \pi^i(\xi) &= \sum_{\zeta \in \mathcal{A}^i} \pi^i(\xi \cap \zeta) \\ &= \mu \sum_{\zeta \in \mathcal{A}^i} \pi^{i*}(\xi \cap \zeta) = \mu \pi^{i*}(\xi), \end{aligned}$$

and we have $\mu = 1/\pi^{i*}(\xi)$. As a consequence, all conditional probabilities of events $\xi \cap \zeta$ are defined as the unique Bayesian updates

$$\pi^i(\xi \cap \zeta) = \frac{\pi^{i*}(\xi \cap \zeta)}{\pi^{i*}(\xi)} = \pi^{i*}(\zeta|\xi).$$

Combining this restriction with Bayes' law leads to the implication

$$\pi^i \in \Pi^i(\xi) \implies \exists \pi^{i*} \in \arg \max_{\hat{\pi} \in \Pi^i} \hat{\pi}(\xi) \text{ with } \pi^i(\cdot) = \pi^{i*}(\cdot|\xi),$$

and the maximum-likelihood rule (3) follows.

Sufficiency of $\Pi_{\mathbb{M}}^i$: It is obvious that Bayes' law is implied; thus, it remains to be shown that the rule satisfies ambiguity containment. Under the rule, $\pi^i \in \Pi^i(\xi)$ implies

$$\pi^i(\xi \cap \zeta) = \pi^{i*}(\zeta|\xi) = \frac{\pi^{i*}(\xi \cap \zeta)}{\pi^{i*}(\xi)} \quad \forall \zeta \in \mathcal{A}^i,$$

and thus it holds for the likelihood ratio of any two events $\zeta, \zeta' \in \mathcal{A}^i$ that

$$\frac{\pi^i(\xi \cap \zeta)}{\pi^i(\xi \cap \zeta')} = \frac{\pi^{i*}(\xi \cap \zeta)}{\pi^{i*}(\xi \cap \zeta')} = \frac{\max_{\hat{\pi} \in \Pi^i} \hat{\pi}(\xi \cap \zeta)}{\max_{\hat{\pi} \in \Pi^i} \hat{\pi}(\xi \cap \zeta')}.$$

The second equality follows from Lemma 2 and establishes ambiguity containment. \square

Proof of Proposition 2. The equivalence is proven by establishing two separate implications.

(i) *implies* (ii): Let $(\xi, \zeta) \in \mathcal{I}_t \times \mathcal{A}^i$ for some $t > 0$. As in Equation (3), let $\pi^{i*} \in \arg \max_{\pi^i \in \Pi^i} \pi^i(\xi)$. First, it shall be shown that

$$\xi \cap \zeta = \xi \implies \pi^i(\xi \cap \zeta|\xi) = \pi^{i*}(\xi \cap \zeta|\xi) \quad \forall \pi^i \in \Pi^i. \quad (21)$$

Since $\xi \cap \zeta = \xi$ means $\xi \subseteq \zeta$, it follows that $\pi^i(\xi \cap \zeta|\xi) = \pi^i(\xi|\xi) = 1 \quad \forall \pi^i \in \Pi^i$ and (21) is trivially true. Second, it shall be shown that

$$\xi \cap \zeta = \zeta \implies \pi^i(\xi \cap \zeta|\xi) = \pi^{i*}(\xi \cap \zeta|\xi) \quad \forall \pi^i \in \Pi^i. \quad (22)$$

Since $\xi \cap \zeta = \zeta$ means that $\xi \cap \zeta \in \mathcal{A}^i$, there is no ambiguity about its probability, i.e., $\pi^i(\xi \cap \zeta) = \pi^{i*}(\xi \cap \zeta) \quad \forall \pi^i \in \Pi^i$. As all unconditional probabilities agree, so must the conditional. Third, note that

$$\xi \cap \zeta = \emptyset \implies \pi^i(\xi \cap \zeta|\xi) = 0 \quad \forall \pi^i \in \Pi^i. \quad (23)$$

Finally, since (21), (22), and (23) hold for any $\xi \in \mathcal{I}_t, t > 0$, there can be no $\pi^i \in \Pi_{\mathbb{B}}^i(\xi)$ such that $\pi^i \notin \Pi_{\mathbb{M}}^i(\xi)$. As $\Pi_{\mathbb{M}}^i(\xi) \subseteq \Pi_{\mathbb{B}}^i(\xi)$ by definition, the implication is proven.

$\neg(i)$ *implies* $\neg(ii)$: It shall be shown that for any $(\xi, \zeta) \in \mathcal{I}_t \times \mathcal{A}^i, t > 0$,

$$\xi \cap \zeta \notin \{\xi, \zeta, \emptyset\} \implies \exists \pi^i \in \Pi^i \text{ such that } \pi^i(\xi) \neq \pi^{i*}(\xi). \quad (24)$$

Since $\xi \cap \zeta \neq \zeta$ and $\xi \cap \zeta \neq \emptyset$ jointly imply that $\xi \cap \zeta \notin \mathcal{A}^i$, there exists some $\pi^i \in \Pi^i$ such that $\pi^i(\xi \cap \zeta) \neq \pi^{i*}(\xi \cap \zeta)$; otherwise, \mathcal{A}^i would not be the finest partition satisfying (4). By Lemma 2,

$$\pi^{i*}(\xi \cap \zeta) = \max_{\hat{\pi}^i \in \Pi^i} \hat{\pi}^i(\xi \cap \zeta)$$

for any $\zeta' \in \mathcal{A}^i$. Therefore, for at least one cell $\zeta \in \mathcal{A}^i$, $\pi^i(\xi \cap \zeta) < \pi^{i*}(\xi \cap \zeta)$ and for all other cells $\zeta' \neq \zeta$, $\pi^i(\xi \cap \zeta') \leq \pi^{i*}(\xi \cap \zeta')$. Summing over all cells of \mathcal{A}^i leads to

$$\sum_{\zeta' \in \mathcal{A}^i} \pi^i(\xi \cap \zeta') = \pi^i(\xi) < \pi^{i*}(\xi),$$

and thus (24) holds true. This proves the implication and thus the proposition. \square

Proof of Proposition 3. For any set $S \subset \mathbb{R}^{|\Omega|}$, define the rescale operation $R(S) = \left\{ \frac{x}{x \cdot \vec{1}} \mid x \in S \right\}$ and the maximizer $M_\xi(S) = \arg \max_{x \in S} x \cdot \vec{1}$. Then, full Bayesian updating and maximum-likelihood updating can be written as

$$\Pi_{\mathbb{B}}^i(\xi) = R(\text{pr}_\xi(\Pi^i)) \quad (25)$$

$$\Pi_{\mathbb{M}}^i(\xi) = R(\text{pr}_\xi(M_\xi(\Pi^i))). \quad (26)$$

We show that for any $t > 0$, a recursive application of the rectangularity condition (Definition 9) from date 0 through date $t - 1$ implies that the two updating rules agree on all date t events: At date 0, $\xi_0 = \Omega$ for all $\xi_0 \in \mathcal{I}_0$, and thus $\Pi_{\mathbb{B}}^i(\xi_0) = \Pi^i$. Using rectangularity recursively leads to

$$\Pi^i = \bigcup_{\pi_1^i \in \Pi^i} \sum_{\xi_1 \in \mathcal{I}_1} \pi_1^i(\xi_1) \bigcup_{\pi_2^i \in \Pi^i} \sum_{\xi_2 \in \mathcal{I}_2} \pi_2^i(\xi_2 | \xi_1) \cdots \bigcup_{\pi_t^i \in \Pi^i} \sum_{\xi_t \in \mathcal{I}_t} \pi_t^i(\xi_t | \xi_{t-1}) \Pi_{\mathbb{B}}^i(\xi_t). \quad (27)$$

For any $\xi \in \mathcal{I}_t$, the maximum operation M_ξ can be applied to both sides of (27). Let $\hat{\xi}_0, \hat{\xi}_1, \dots, \hat{\xi}_t$ be the sequence of events $\hat{\xi}_s \in \mathcal{I}_s$ that satisfies $\hat{\xi}_t = \xi$ and $\hat{\xi}_s \subseteq \hat{\xi}_{s-1}$ for all $1 \leq s \leq t$. Note that M_ξ applied to the right-hand side of (27) maximizes the product $\pi_1^i(\hat{\xi}_1) \pi_2^i(\hat{\xi}_2 | \hat{\xi}_1) \cdots \pi_t^i(\hat{\xi}_t | \hat{\xi}_{t-1})$. Since the constraint set $(\Pi^i)^t = \times_{s=1}^t \Pi^i$ has a product structure,

$$\max_{(\pi_1^i, \dots, \pi_t^i) \in (\Pi^i)^t} \prod_{s=1}^t \pi_s^i(\hat{\xi}_s | \hat{\xi}_{s-1}) = \prod_{s=1}^t \max_{\pi_s^i \in \Pi^i} \pi_s^i(\hat{\xi}_s | \hat{\xi}_{s-1}) = \lambda^*,$$

in which λ^* is a strictly positive scalar because each maximum must be greater than zero under Assumption 1. Using M_ξ on both sides of (27) thus leads to

$$M_\xi(\Pi^i) = \bigcup_{\pi_1^i \in M_{\hat{\xi}_1}(\Pi^i)} \sum_{\xi_1 \in \mathcal{I}_1} \pi_1^i(\xi_1) \cdots \bigcup_{\pi_t^i \in M_{\hat{\xi}_t}(\Pi^i)} \sum_{\xi_t \in \mathcal{I}_t} \pi_t^i(\xi_t | \xi_{t-1}) \Pi_{\mathbb{B}}^i(\xi_t). \quad (28)$$

Note that $\text{pr}_\xi(\Pi_{\mathbb{B}}^i(\xi_t)) = \{0\}$ for all $\xi_t \in \mathcal{I}_t \setminus \{\xi\}$. Therefore, pr_ξ applied to both sides of (28) results in

$$\text{pr}_\xi(M_\xi(\Pi^i)) = \bigcup_{\pi_1^i \in M_{\hat{\xi}_1}(\Pi^i)} \cdots \bigcup_{\pi_t^i \in M_{\hat{\xi}_t}(\Pi^i)} \text{pr}_\xi \left(\pi_1^i(\hat{\xi}_1) \pi_2^i(\hat{\xi}_2 | \hat{\xi}_1) \cdots \pi_t^i(\hat{\xi}_t | \hat{\xi}_{t-1}) \Pi_{\mathbb{B}}^i(\xi) \right), \quad (29)$$

and the rescale operation R applied to both sides of (29) in

$$R(\text{pr}_\xi(M_\xi(\Pi^i))) = R(\lambda^* \text{pr}_\xi(\Pi_{\mathbb{B}}^i(\xi))). \quad (30)$$

By construction, $R(\lambda S) = R(S)$ for any scalar $\lambda \neq 0$ and set S . Therefore, by means of (25) and (26), Equation (30) can be rewritten as

$$\Pi_{\mathbf{M}}^i(\xi) = \Pi_{\mathbf{B}}^i(\xi). \quad (31)$$

Since this is true for any $t > 0$ and any choice of $\xi \in \mathcal{I}_t$, the proposition is proven. \square

The proof of Corollary 3 employed the following Lemma.

Lemma 3. *If the financial market is dynamically complete at a dynamically consistent equilibrium (P, C, Ψ) , then $Q[P_0]$ is a one-dimensional subspace.*

Proof. Since $Q[P_0]$ is the orthogonal complement $Q[P_0] = X(P_0, \mathcal{S}^J; \Omega)^\perp$ of a linear subspace, it is a subspace itself and satisfies

$$\dim(Q[P_0]) + \dim(X(P_0, \mathcal{S}^J; \Omega)) = \dim(\mathcal{S}) = \sum_{t=0}^T |\mathcal{I}_t|. \quad (32)$$

Suppose $\dim(Q[P_0]) = 0$; then, $X(P_0, \mathcal{S}^J; \Omega) = \mathcal{S}$ and the budget set $B_0^i(P_0, 0)$ is not compact. In this case, no undominated plan and thus no equilibrium exists. Therefore, at any equilibrium

$$\dim(Q[P_0]) \geq 1. \quad (33)$$

Note that $X_{t+1}(P_0, \mathcal{S}^J; \xi)$ is a subspace of $X(P_0, \mathcal{S}^J; \Omega)$ for any $\xi \in \mathcal{I}_t$ and $t \in \{0, \dots, T-1\}$. Since \mathcal{I}_t is a partition, $X_{t+1}(P_0, \mathcal{S}^J; \xi)$ and $X_{t+1}(P_0, \mathcal{S}^J; \xi')$ are transverse for any $\xi, \xi' \in \mathcal{I}_t$ such that $\xi \neq \xi'$. Moreover, $X_s(P_0, \mathcal{S}^J; \cdot)$ and $X_t(P_0, \mathcal{S}^J; \cdot)$ are transverse for any $s \neq t$. This leads to

$$\begin{aligned} \dim \left(\sum_{t=0}^{T-1} \sum_{\xi \in \mathcal{I}_t} X_{t+1}(P_0, \mathcal{S}^J; \xi) \right) &= \sum_{t=0}^{T-1} \sum_{\xi \in \mathcal{I}_t} \dim(X_{t+1}(P_0, \mathcal{S}^J; \xi)) \\ &\leq \dim(X(P_0, \mathcal{S}^J; \Omega)), \end{aligned} \quad (34)$$

in which the equality follows from transversality while the inequality follows from the fact that any sum of subspaces is contained in the whole space. Dynamic consistency implies that $P_0 = \dots = P_T$, and thus Definition 6 implies that for each $t \geq 0$,

$$\sum_{\xi \in \mathcal{I}_t} \dim(X_{t+1}(P_0, \mathcal{S}^J; \xi)) = |\mathcal{I}_{t+1}|$$

because \mathcal{I}_t is a partition. Inserting this result into (34) leads to

$$\sum_{t=1}^T |\mathcal{I}_t| \leq \dim(X(P_0, \mathcal{S}^J; \Omega)),$$

and, in combination with (32), this results in

$$\dim(Q[P_0]) \leq |\mathcal{I}_0| = 1. \quad (35)$$

Jointly, (33) and (35) imply that $\dim(Q[P_0]) = 1$, which concludes the proof of the lemma. \square

References

- Arrow, K. J., and Debreu, G. 1954. Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22 (8), 265–290.
- Bewley, T. F. 1986. Knightian Decision Theory. Part I. Discussion Paper No. 807, Cowles Foundation.
- Bewley, T. F. 2002. Knightian Decision Theory. Part I. *Decisions in Economics and Finance* 25 (2), 79–110.
- Citanna, A., Kajii, A., and Villanacci, A. 1998. Constrained Suboptimality in Incomplete Markets: A General Approach and Two Applications. *Economic Theory* 11 (3), 495–521.
- Cohen, M., Gilboa, I., Jaffray, J.-Y., and Schmeidler, D. 2000. An Experimental Study of Updating Ambiguous Beliefs. *Risk, Decision, and Policy* 5 (2), 123–133.
- Dana, R.-A., and Riedel, F. 2013. Intertemporal Equilibria with Knightian Uncertainty. *Journal of Economic Theory* 148 (4), 1582–1605.
- Dominiak, A., Duersch, P., and Lefort, J.-P. 2012. A Dynamic Ellsberg Urn Experiment. *Games and Economic Behavior* 75 (2), 625–638.
- Epstein, L. G. 1999. A Definition of Uncertainty Aversion. *Review of Economic Studies* 66 (3), 579–608.
- Epstein, L. G., and Schneider, M. 2003. Recursive Multiple-Priors. *Journal of Economic Theory* 113 (1), 1–31.
- Geanakoplos, J., and Polemarchakis, H. M. 1986. Existence, Regularity, and Constrained Suboptimality of Competitive Allocations when the Asset Market is Incomplete. In W. P. Heller, R. M. Starr, and D. A. Starrett(ed.) *Uncertainty, Information and Communication: Essays in Honor of Kenneth J. Arrow, Volume III*: Cambridge Cambridge University Press.
- Gerasimou, G. 2018. On the Indifference Relation in Bewley Preferences. *Economics Letters* 164, 24–26.
- Gilboa, I., Maccheroni, F., Marinacci, M., and Schmeidler, D. 2010. Objective and Subjective Rationality in a Multiple Prior Model. *Econometrica* 78 (2), 755–770.
- Gilboa, I., and Schmeidler, D. 1989. Maxmin Expected Utility with Non-unique Prior. *Journal of Mathematical Economics* 18 (2), 141–153.
- Gilboa, I., and Schmeidler, D. 1993. Updating Ambiguous Beliefs. *Journal of Economic Theory* 59 (1), 33–49.

- Hanany, E., and Kilbanoff, P. 2007. Updating Preferences with Multiple Priors. *Theoretical Economics* 2 (3), 261–298.
- Herings, P. J.-J., and Rohde, K. I. M. 2008. On the Completeness of Complete Markets. *Economic Theory* 37 (2), 171–201.
- Knight, F. H. 1921. *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.
- Kreps, D. M. 1982. Multiperiod Securities and the Efficient Allocation of Risk: A Comment on the Black-Scholes Option Pricing Model. In J. J. McCall(ed.) *The Economics of Information and Uncertainty*: Chicago University of Chicago Press.
- Lucas, Jr., R. E., and Prescott, E. C. 1971. Uncertainty and Risk in Financial Markets. *Econometrica* 39 (5), 659–681.
- Ma, W. 2015. The Existence and Efficiency of General Equilibrium with Incomplete Markets. *Economics Letters* 134, 78–81.
- Ma, W. 2018. Real Indeterminacy of General Equilibrium under Knightian Uncertainty. *Journal of Mathematical Economics* 79, 106–111.
- Magill, M., and Quinzii, M. 1996. *Theory of Incomplete Markets*. Cambridge, London: MIT Press.
- Magill, M. J., and Shafer, W. J. 1990. Characterisation of Generically Complete Real Asset Structures. *Journal of Mathematical Economics* 19 (1-2), 167–194.
- Mandler, M. 2013. Endogenous Indeterminacy and Volatility of Asset Prices under Ambiguity. *Theoretical Economics* 8 (3), 729–750.
- Mukerji, S., and Tallon, J.-M. 2001. Ambiguity Aversion and Incompleteness of Financial Markets. *Review of Economic Studies* 68 (4), 883–904.
- Ngangoué, K. 2018. Learning under Ambiguity: An Experiment on Gradual Information Processing. working paper.
- Radner, R. 1970. Problems in the Theory of Markets under Uncertainty. *American Economic Review* 60 (2), 454–460.
- Radner, R. 1972. Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets. *Econometrica* 40 (2), 289–303.
- Rigotti, L., and Shannon, C. 2005. Uncertainty and Risk in Financial Markets. *Econometrica* 73 (1), 203–243.
- Rigotti, L., and Shannon, C. 2012. Sharing Risk and Ambiguity. *Journal of Economic Theory* 147 (5), 2028–2039.
- Savage, L. J. 1954. *The Foundations of Statistics*. New York: Wiley.