

Dynamic Inconsistency and Inefficiency of Equilibrium under Knightian Uncertainty*

Patrick Beissner[†] Michael Zierhut[‡]

June 19, 2022

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 outcomes varies with choice of updating rule. We document two phenomena: First, unless agents apply the full Bayesian rule, consumption decisions may be dynamically inconsistent. Second, unless they apply the maximum-likelihood rule, ambiguous probability mass may dilate, which causes price fluctuations. Either phenomenon results in Pareto inefficient allocations. We ask whether it is possible to design one updating rule that prevent both phenomena. The answer is negative: No such rule exists. Efficiency can be restored by restricting priors: Full Bayesian and maximum-likelihood updates agree when priors are rectangular, and when ambiguity is sufficiently large, all equilibria are Pareto efficient, even if prices and allocations change over time.

JEL Classification: D52; D53; D81

Keywords: Knightian uncertainty, Dynamic consistency, Dilation, Maximum rectangularity, General equilibrium with incomplete preferences

*The authors thank Sophie Bade, Martin Barbie, Pablo Beker, Federico Echenique, Georgios Gerasi-mou, Chiaki Hara, Christian Kellner, Fabio Maccheroni, Filipe Martins-da-Rocha, Even Piermont, Frank Riedel, Marciano Siniscalchi, Ronald Stauber, Norio Takeoka, Takashi Ui, Georg Weizsäcker, Jan Werner, Ryoko Wada, Xinyang Wang, Łukasz Woźny, Horst Zank, as well as participants in the workshop on Mathematics of Behavioral Economics and Knightian Uncertainty in Financial Markets at ZiF Bielefeld, the 2018 North American Summer Meeting of the Econometric Society, the conference on Time, Uncertainties and Strategies VI, the 2021 SAET Conference, and in seminars at Hitotsubashi University, Kyoto University, Instituto Tecnológico Autónomo de México, University of Cologne, and Royal Holloway, University of London for helpful comments. M. Zierhut gratefully acknowledges financial support by the German Research Foundation (DFG) under projects ZI 1673/1-1 and ZI 1673/3-1.

[†]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 preferences of the Bewley type meet in a sequence of complete contingent markets à la Arrow and Debreu (1954).

We build on the important contribution of Rigotti and Shannon (2005), who generalize the two-date Arrow-Debreu model to a setting with multiple priors. Two of their findings are particularly relevant in our context: First, Knightian uncertainty results in a large equilibrium set, including equilibria with identical prices but different allocations, as well as equilibria with identical allocations but different prices. Second, all these equilibria are Pareto efficient. This welfare result is strong but limited to a market that opens 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 modifies this setting by permitting trade at subsequent dates. Under subjective expected utility, this modification would have no effect: Agents attain a Pareto efficient allocation in the initial market, and even if markets open again in the future, no more trade takes place because preference orderings and budget sets do not change. This turns out to be different under Knightian uncertainty.

In contrast with the two-date model, beliefs must be updated on several consecutive dates. From studies of the maxmin expected utility model of Gilboa and Schmeidler (1989), it is known that updates of set-valued beliefs may reverse preference orderings. Such dynamic inconsistency also occurs in the Knightian decision model, and agents may revise their initial consumption plans, even if their budget sets do not change. A different phenomenon at play is what Seidenfeld and Wasserman (1993) call *dilation*: In response to new information, ambiguous probability mass may spread to unambiguous events. Agents become uncertain about probabilities they knew before. In the equilibrium context, dilation results in a sudden expansion of the set of market-clearing prices. This gives rise to equilibria at which budget sets change over time, and agents revise their initial consumption plans, even if their preferences are dynamically consistent. Both

phenomena result in deviations from Pareto efficient allocations after new information is released.

Whether dynamic inconsistency and dilation occur depends on the choice of updating rule for beliefs. In the traditional single-prior setting, Bayes' law determines a unique updating rule. In settings 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. As Hanany and Kilbanoff (2007) show, none of these rules are able to guarantee dynamic consistency in the maxmin expected utility model. In the Knightian decision model, a more affirmative result can be obtained: The criterion of dynamic consistency is fulfilled if agents update their beliefs prior by prior, as suggested by Bewley (1987).

1. The unique updating rule that guarantees dynamic consistency is full Bayesian updating. (Proposition 2)

Dilation goes against the principle that better information should be used to reduced uncertainty. Contrary to intuition, agents forget probabilities they knew before. We define a condition of ambiguity containment, which ensures that new information does not create new uncertainty. This condition is met if agents apply the maximum-likelihood rule, as axiomatized by Gilboa and Schmeidler (1993).

2. The unique updating rule that guarantees ambiguity containment is maximum-likelihood updating. (Proposition 3).

Since both these rationality criteria are standard in the single-prior setting, a natural objective is to identify an updating rule that guarantees both jointly, even if agents have multiple priors. However, this objective cannot be achieved: We show that there exists no updating rule that satisfies Bayes' law and guarantees both dynamic consistency and ambiguity containment for arbitrary priors. In light of this impossibility result, we ask whether there is a nontrivial subset of economies in which this conflict between rationality criteria disappears. This question is answered in the affirmative: The rectangularity condition of Epstein and Schneider (2003), which characterizes dynamic consistency in the maximum expected utility model, turns out to be relevant for sequential decisions under Bewley preferences. We offer the following characterization:

3. Rectangularity holds if and only if full Bayesian updating, maximum-likelihood updating, and minimum-likelihood updating agree on observable events (Theorem 1).

Together with the two previous characterizations, this result establishes that rectangularity implies dynamic consistency and ambiguity containment simultaneously. These properties of decision making translate into properties of market equilibrium. It is well understood that incomplete preferences fall short of assumptions on how decisions between incomparable alternatives are made. In absence of such assumptions, agents may switch back and forth between incomparable plans over time. This results in equilibria

with fluctuations in prices and allocations that may be hard to justify. To avoid equilibria with erratic behavior, it has become customary to apply equilibrium refinements. To this end, Bewley (1986) proposes the criterion of *inertia*: Plans never change without a reason. Agents will only deviate from the status quo if a preferred alternative becomes available.¹ The combination of dynamic consistency and inertia leads to market outcomes that are socially desirable:

4. Dynamic consistency and inertia jointly imply Pareto efficiency (Theorem 2).

In spite of its normative appeal, inertia is a controversial criterion. It may be criticized for lacking an axiomatic foundation. Some critics may even argue that the volatility it eliminates is a natural feature of Knightian uncertainty. We ask whether it is possible to reach the same welfare standard without resorting to inertia. Again we can answer this question in the affirmative, albeit under strong restrictions on priors. We introduce a condition of *maximum rectangularity*, which combines rectangularity with a maximum of ambiguity about probabilities of events observable at intermediate dates. Contrary to inertia, the idea is not to prevent deviations from a Pareto efficient initial market outcome, but to render the set of Pareto efficient allocations sufficiently large that all deviations remain within this set. This requires that ambiguity does not spread to unambiguous events. In this case, even if prices and allocations change over time, market outcomes are socially desirable:

5. Ambiguity containment and maximum rectangularity jointly imply Pareto efficiency (Theorem 3).

This is the attractive feature of the Knightian decision model: Even though it shares with other incomplete preference models the difficulty that dynamically efficient behavior depends on intertemporal consistency conditions, these conditions boil down to properties of updating rules. These properties are testable and have a normative foundation in decision theory.

The remainder of this paper is structured as follows. Section 2 introduces the model. Section 3 demonstrates, by means of an example, puzzling features of updating rules and the resulting inefficiency. Section 4 connects updating rules with normative criteria, and presents all decision theoretic results. This section is self-contained and accessible after preliminary reading of sections 2.1 through 2.3. Section 5 focuses on equilibrium aspects and presents all welfare results. Section 6 concludes.

2 Model

Consider a finite-horizon stochastic economy with symmetric information. The stochastic model is a tree: The terminal nodes represent possible states of the world, and each

¹Concerned with but a single agent, Bewley speaks of inertia as a *behavioral assumption*, not as part of a solution concept.

predecessor node represents public information at a given date about the realized state. Information is gradually released over a finite number of dates. The economy is populated by a finite number of agents whose objects of choice are contingent consumption plans: Each agent plans how much to consume at each node of the tree. Contingent markets can be used to transfer consumption across different nodes.

2.1 Time and Uncertainty

There are $T + 1$ dates, enumerated by $t \in \{0, \dots, T\}$. At date 0, the state of the world is drawn from a finite state space Ω but not revealed. At each date $t > 0$, new information about the state of the world becomes public. This gradual revelation is modeled as a sequence of information partitions $\mathcal{I} = \{\mathcal{I}_t\}_{t=0}^T$ on Ω . 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 . 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 \mathcal{I} defines a tree of *date-events* (t, ξ) , with $L = \sum_{t=0}^T |\mathcal{I}_t|$ nodes.

Probabilities on Ω are represented by a measure $\mu : 2^\Omega \rightarrow [0, 1]$ normalized to $\mu(\Omega) = 1$. Since the number of states is finite, every probability measure can be represented as a point on $\Delta^{|\Omega|}$, the unit simplex in $|\Omega|$ -dimensional Euclidean space. A state $\omega \in \Omega$ occurs with probability $\mu(\omega)$, and an event $\xi \subseteq \Omega$ occurs with probability $\mu(\xi)$. For any subset of events $\mathcal{E} \subset 2^\Omega$, denote by $\mu|_{\mathcal{E}}$ the domain restriction of μ to \mathcal{E} . Moreover, denote by $\mathbb{E}^\mu[\cdot]$ the expectation operator under μ .

2.2 Agents and Preferences

There are I agents, indexed by superscripts $i \in \{1, \dots, I\}$. All agents have preferences over consumption plans from a common consumption space $\mathcal{C} = \mathbb{R}^L$. There is one commodity per date-event, and each consumption plan $x \in \mathcal{C}$ maps date-events (t, ξ) to consumption levels $x(t, \xi)$. The following notation is used throughout: $x \geq 0$ 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. The positive orthant is $\mathcal{C}_+ = \{x \in \mathcal{C} \mid x \geq 0\}$, and its interior is $\mathcal{C}_{++} = \{x \in \mathcal{C} \mid x \gg 0\}$.

The preference relation $x \succ^i y$ indicates that agent i prefers $x \in \mathcal{C}_+$ to $y \in \mathcal{C}_+$. All agents have preferences of the Bewley type: Each agent i has multiple priors, represented by a set Π^i of subjective probabilities $\pi^i : 2^\Omega \rightarrow [0, 1]$.

Assumption 1 (Subjective probabilities). For each agent i , Π^i is closed and convex. Each event $\xi \in \mathcal{I}_t$, $t \geq 0$, has a positive probability $\pi^i(\xi) > 0$ under each $\pi^i \in \Pi^i$. Each state $\omega \in \Omega$ has a positive probability $\pi^i(\omega) > 0$ under some $\pi^i \in \Pi^i$.

The positivity requirements in Assumption 1 are mild regularity conditions. These

rule out division by zero when computing conditional probabilities as well as degenerate priors that are entirely contained in one face of $\Delta^{|\Omega|}$.

Assumption 2 (Preferences). For each agent i ,

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

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

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.² In this case, the preference relation can be represented by a collection of implied expected utility functions $U^{\pi^i}(x) = \mathbb{E}^{\pi^i}[u^i(x)]$, one for each $\pi^i \in \Pi^i$. A plan x is said to be *undominated* if there is some $\pi^i \in \Pi^i$ such that x maximizes U^{π^i} over the set of alternatives; this means, there is no feasible alternative $y \succ^i x$. Some plans are always comparable: If $x \gg y$, then the expected utility $U^{\pi^i}(x)$ must exceed the expected utility $U^{\pi^i}(y)$ under any probability measure π^i . By differentiability, all implied expected utility functions have a well-defined gradient $DU^{\pi^i}[x] \in \mathcal{C}_{++}$ at any interior consumption plan x . These gradients can be used to define the *marginal rates cone* of the preference relation \succ^i as the conical hull

$$\nabla U^i[x] = \text{cone} \left\{ DU^{\pi^i}[x] \mid \pi^i \in \Pi^i \right\}.$$

Bewley's inertia assumption requires that agents compare alternatives to some status quo. The status quo of agent i at date 0 is his endowment $e^i \in \mathcal{C}_{++}$. If agents decide at a future date $t > 0$ to deviate from a status quo, choices are restricted to the conditional domain

$$\mathcal{C}_t(e^i) = \{x \in \mathcal{C}_+ \mid x(s, \xi') = e^i(s, \xi') \quad \forall \xi' \in \mathcal{I}_s \quad \forall s < t\};$$

that is to say, only future consumption can be changed and the past cannot be undone. Further constraints for choices at a specific date-event (t, ξ) will be formulated using the *subtree inner product*

$$[x, y]_{(t, \xi)} = \sum_{s=0}^T \sum_{\xi' \subseteq \xi} x(s, \xi') y(s, \xi') \mathbb{1}_{\mathcal{I}_s}(\xi'),$$

which agrees with the canonical dot inner product $x \cdot y$ at the root node $(0, \Omega)$.

²This preference representation is introduced by Rigotti and Shannon (2005) and is slightly more general than the axiomatization of Bewley (1986, 2002), which results in a linear function u^i . Additive separability is a further generalization of Bewley (1987), who assumes a function $u^i(x) = \sum_{t=0}^T \beta^t \tilde{u}^i(x_t)$ with constant discount factor $\beta \in (0, 1)$.

2.3 Updating

To make precise the idea that agents update their subjective probabilities in light of new information, define the *updating rule* $\Pi^i : 2^\Omega \rightrightarrows \Delta^{|\Omega|}$ of agent i as a nonempty-closed-convex-valued correspondence from the observed event $\xi \subseteq \Omega$ to his set of posterior probabilities $\Pi^i(\xi)$. We follow the convention of dropping parentheses whenever Ω is the argument, such that $\Pi^i(\Omega) = \Pi^i$ is 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 multiple-prior case, the set of candidate updating rules is large. 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. The conditional preference relation of agent i on observing ξ is jointly determined by Assumption 2 and the update $\Pi^i(\xi)$:

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

In words: x is conditionally preferred to y in event ξ . Conditional preferences result in a conditional marginal rates cone

$$\nabla U^i[x](\xi) = \text{cone} \left\{ DU^{\pi^i}[x] \mid \pi^i \in \Pi^i(\xi) \right\}.$$

We shall write $x \succ_t^i y$ if x is conditionally preferred to y in some event $\xi \in \mathcal{I}_t$ and conditionally equivalent or preferred to y in each event $\xi \in \mathcal{I}_t$, in which conditional equivalence means that x and y have identical values in the entire subtree rooted at date-event (t, ξ) . As we shall see, two candidates for the updating rule Π^i are particularly relevant from a normative viewpoint. The first is the so-called *full Bayesian rule*:

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

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*:

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

One additional updating rule that has little normative appeal but interesting technical properties is the *minimum-likelihood rule*, which reduces model uncertainty by discarding all models except those least likely to generate the observed data:³

$$\Pi_{\mathbb{W}}^i(\xi) = \left\{ \pi^{i^*}(\cdot | \xi) \mid \pi^{i^*} \in \arg \min_{\pi^i \in \Pi^i} \pi^i(\xi) \right\}. \quad (4)$$

2.4 Production

There are K production technologies indexed by superscripts $k \in \{1, \dots, K\}$. Each technology can be used to convert some consumption good at date t to contingent consumption at date $t + 1$. All agents have non-rivalrous access to the same K technologies. The key assumption is that production plans are short-lived: If production starts at date-event (t, ξ) , the costs of production have to be paid immediately, and all output becomes available at the direct successor nodes in the date-event tree. Once a successor node is reached, production plans for the next period are implemented, and this continues until the final output is paid off at the terminal date T . Production plans entail no long-term commitment.

All feasible production plans are collected in the production set $\mathcal{Y}^k \subset \mathcal{C}$. For each element $Y \in \mathcal{Y}^k$, negative signs $Y(t, \omega) < 0$ denote input and positive signs $Y(t, \omega) > 0$ denote output.

Assumption 3 (Production technology). For each technology k , \mathcal{Y}^k is a closed, convex cone, and there is a date-event (t, ξ) with $t < T$ such that each $Y \in \mathcal{Y}^k$ satisfies

1. $Y \in \mathcal{C}_t(0)$ (no retroaction)
2. $Y(t, \xi) \leq 0$ (irreversibility)
3. $Y(t, \xi) = 0 \implies Y = 0$ (no free output)
4. $Y(s, \xi') > 0 \implies s = t + 1$ (one-period activity)
5. $Y(t + 1, \xi') > 0 \implies \xi' \in \mathcal{I}_{t+1}$ and $\xi' \subseteq \xi$ (output in successor nodes).

Since production sets are cones under Assumption 3, all technologies have constant returns to scale. The following requirements are expressed by the five conditions in Assumption 3: First, there is no input or output before the initial date-event. Second, production requires input at the initial date-event. Third, zero input leads to zero output. Fourth, there is no long-term commitment. Fifth, output cannot appear in other branches of the date-event tree. The aggregate production set is $\bar{\mathcal{Y}} = \sum_{k=1}^K \mathcal{Y}^k$, and if $\bar{\mathcal{Y}} = \{0\}$, the model boils down to a pure exchange economy.

³The subscript is an M turned upside down, just like this rule turns the logic of maximum-likelihood updating upside down.

2.5 Contingent Markets

Trade takes place in a sequence of complete markets for contingent consumption. At date 0, consumption at all L date-events is traded, and all L prices are collected in the price vector $q_0 \in \mathcal{C}_+$. Each agent i chooses a consumption plan $c_0^i \in \mathcal{C}_+$ subject to the date 0 budget constraint

$$q_0 \cdot c_0^i \leq q_0 \cdot e^i + q_0 \cdot Y \quad \text{for some } Y \in \bar{\mathcal{Y}}.$$

In words: The market value of consumption must not exceed the market value of the initial status quo allocation e^i net of feasible production. Since preferences are monotone under Assumptions 1 and 2, the budget constraint will always bind and should be thought of as holding with equality. At date 1, markets open again for consumption at all subsequent nodes of the date-event tree. In light of new information, prices may change to $q_1 \in \mathcal{C}_1(q_0)$, and each agent i may retrade the status quo allocation c_0^i for a revised consumption plan $c_1^i \in \mathcal{C}_1(c_0^i)$. Note that c_1^i is a complete contingent plan of the agent, which incorporates the choices in all markets that are open at date 1 nodes. By definition of the conditional domain, the price and allocation of date 0 consumption good can no longer change at this point in time.

This procedure repeats itself: At any date $t < T$, markets open again after new information is released. Prices may change to $q_t \in \mathcal{C}_t(q_{t-1})$, and the status quo allocation of agent i is c_{t-1}^i , the consumption plan carried over from the previous date. The generic date t budget set of the agent is therefore

$$B_t(q_t, c_{t-1}^i) = \left\{ c_t^i \in \mathcal{C}_t(c_{t-1}^i) \mid \begin{array}{l} [q_t, c_t^i - c_{t-1}^i]_{(t,\xi)} = [q_t, Y]_{(t,\xi)} \\ \text{for all } \xi \in \mathcal{I}_t \text{ and some } Y \in \bar{\mathcal{Y}} \end{array} \right\},$$

which depends on current prices and on the current status quo allocation. At date 0, this status quo is understood as $c_{-1}^i = e^i$. The choices of all agents at date t are collected in the vector $c_t = (c_t^1, \dots, c_t^I)$.

2.6 Equilibrium

Agents face a sequence of prices $q = (q_0, \dots, q_{T-1})$ and a sequence of decisions $c = (c_0, \dots, c_{T-1})$. Recall that when preferences are complete, the decision criterion is optimality. A decision is *optimal* if it is preferred to all alternatives. By contrast, Bewley preferences are incomplete and thus based on the weaker criterion of undominatedness. A decision is *undominated* if there is no preferred alternative.⁴ The consumption plan c_t^i is undominated for agent i if $\mathcal{D}_t^i(q_t, c_t^i) = \emptyset$, in which $\mathcal{D}_t^i : \mathcal{C}_+ \times \mathcal{C}_+ \rightrightarrows \mathcal{C}_+$ is the *dominating-set correspondence*, defined as

$$\mathcal{D}_t^i(q_t, x) = \{ c_t^i \in B_t(q_t, x) \mid c_t^i \succ_t^i x \}.$$

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

The economy is in equilibrium if all agents choose undominated consumption plans subject to their budget constraint at given prices, and prices are set in such way that all contingent markets clear.

Definition 2. An *equilibrium* is tuple $(q, c) \in \mathcal{C}_{++}^T \times \mathcal{C}_{++}^{IT}$ of prices and consumption plans that satisfies at each date $t \in \{0, \dots, T-1\}$,

1. for each agent i ,

$$\mathcal{D}_t^i(q_t, c_t^i) = \emptyset \quad \text{and} \quad (c_t^i \neq c_{t-1}^i \implies c_t^i \in B_t(q_t, c_{t-1}^i))$$

2. market clearing,

$$\sum_{i=1}^I (c_t^i - e^i) \in \bar{\mathcal{Y}}.$$

Note that equilibrium prices and allocations are strictly positive by definition: Given the results of Rigotti and Shannon (2005), we expect a large number of equilibria, indeterminate both in prices and allocations. In face of such indeterminacy, we take the liberty of concentrating on interior equilibria, at which utility functions U^{π^i} are differentiable. Equilibrium indeterminacy reduces the explanatory power of the model, and this side-effect of incomplete preferences is particularly pronounced in a dynamic context: As time progresses, agents may switch back and forth between undominated choices for no particular reason. In awareness of this issue, Bewley (1986, 2002) introduces the assumption of *inertia*: An agent deviates from a status quo only if the status quo is dominated by the new choice. Inertia leads to an equilibrium refinement in the allocation dimension. This is what Rigotti and Shannon (2005) and Dana and Riedel (2013) refer to as equilibrium with inertia:

Definition 3. An *equilibrium with inertia* is tuple $(q, c) \in \mathcal{C}_{++}^T \times \mathcal{C}_{++}^{IT}$ of prices and consumption plans that satisfies at each date $t \in \{0, \dots, T-1\}$,

1. for each agent i ,

$$\mathcal{D}_t^i(q_t, c_t^i) = \emptyset \quad \text{and} \quad (c_t^i \neq c_{t-1}^i \implies c_t^i \in \mathcal{D}_t^i(q_t, c_{t-1}^i))$$

2. market clearing,

$$\sum_{i=1}^I (c_t^i - e^i) \in \bar{\mathcal{Y}}.$$

In comparison with Definition 2, the budget constraint $c_t^i \in B_t(q_t, c_{t-1}^i)$ is replaced with the stronger requirement that choices must be affordable *and* dominate the status quo c_{t-1}^i . Recall that the status quo of the agent at date 0 is $c_{-1}^i = e^i$. In both of the above equilibrium concepts, the production set $\bar{\mathcal{Y}}$ places an exogenous restriction on indeterminacy in the price dimension. This insight should be attributed to Chambers

(2014), who characterizes undominated choice in the presence of production technologies. At date 0, an interior consumption plan c_0^i of agent i is undominated if and only if

$$q_0 \in \nabla U^i[c_0^i] \cap N_{\bar{\mathcal{Y}}}[Y] \quad \text{for some } Y \in \bar{\mathcal{Y}}, \quad (5)$$

in which $N_{\bar{\mathcal{Y}}}[Y]$ is the *normal cone* to $\bar{\mathcal{Y}}$ at Y . By convexity of production sets, $N_{\bar{\mathcal{Y}}} = \bigcap_{k=1}^K N_{\mathcal{Y}^k}$ at any point of evaluation. The restrictions imposed by (5) can be summarized as follows: If date-event (t, ξ) is unambiguous for at least one agent i , or contained in the output range of at least one technology k , then there must be unique price $q_0(t, \xi)$ of consumption at that date-event, for any initial equilibrium allocation c_0 . Initial allocations are always *Pareto efficient*: There exists no other feasible allocation \hat{c}_0 such that $\hat{c}_0^i \succ^i c_0^i$ for each agent i . For interior allocations, this follows from the characterization of Pareto efficiency (see Chambers (2014), p. 50, Proposition 6),

$$\bigcap_{i=1}^I \nabla U^i[c_0^i] \cap N_{\bar{\mathcal{Y}}}[Y] \neq \emptyset \quad \text{for some } Y \in \bar{\mathcal{Y}}. \quad (6)$$

Since the first-order condition (5) holds for each agent i in equilibrium, (6) is implied. Therefore, inefficient allocations can only arise if agents abandon their initial consumption plans at a future date. As the following example shows, such inefficient deviations can indeed occur when markets reopen. Technological price restrictions via (5) are used to keep the example tractable.

3 Example

The example is based on 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 $\mathcal{I}_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}\}$, and these seven nodes are assigned labels $\ell = 0, \dots, 6$. We write $\xi_1 = \{\omega_1, \omega_2\}$ and $\xi_2 = \{\omega_3, \omega_4\}$ for the two branches at date 1, corresponding to labels 1 and 2. There are $I = 2$ agents with identical utility functions

$$\mathbb{E}^{\pi^i}[u^i(c^i)] = c^i(0) + \sum_{\ell=1}^L \pi^i(\ell) \ln(c^i(\ell)),$$

and identical endowments $e^i = (6, 0, 0, 0, 1, 1, 0)$. Thus, each agent is endowed with six units of the consumption good at date 0, and one unit at the terminal nodes 4 and 5. Consumption at the remaining four date-events is possible through production. There are $K = 4$ constant returns to scale production technologies that can be used generate output in nodes 1, 2, 3, and 6. One unit of input at the predecessor node yields two units at the output node. These production possibilities are summarized in the aggregate production set $\bar{\mathcal{Y}} = \{\alpha \cdot M \mid \alpha \in \mathbb{R}_+^4\}$, which contains the rows of




Figure 1: Date-event tree. Nodes are labeled with $\ell = 0, \dots, 6$; branches with Agent 1's prior probability of passing through. Relative prices are fixed along solid branches, and endogenous along dashed branches.

$$M = \begin{pmatrix} -1 & 2 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 2 \end{pmatrix}$$

and all their conical combinations. According to the first-order condition (5), equilibrium prices must satisfy the normality condition $q_0 \cdot M^\top = 0$. For example, the first row of M restricts the price of consumption in node 1 relative to node 0 to $\frac{q_0(1)}{q_0(0)} = \frac{1}{2}$. There can be no indeterminacy in these prices. Only the prices of consumption in nodes 4 and 5 may float freely. The corresponding date-event tree is visualized in Figure 1. The only difference between the two agents is in their subjective probabilities. Agent 1 can assign precise probabilities to states ω_2 and ω_3 but faces Knightian uncertainty about the probability of the remaining two states:

$$\Pi^1 = \left\{ \left(\frac{\eta}{2}, \frac{1}{4}, \frac{1}{4}, \frac{1-\eta}{2} \right) \mid \eta \in [0, 1] \right\}$$

From the perspective of Agent 1, the set of elementary unambiguous events is $\mathcal{A}^1 = \{\{\omega_2\}, \{\omega_3\}, \{\omega_1, \omega_4\}\}$. This set is a partition of the state space Ω , and will be referred to as *ambiguity partition*. By contrast, Agent 2 assigns precise probabilities to all four states and is therefore a classical expected utility maximizer:

$$\Pi^2 = \left\{ \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4} \right) \right\}$$

The set of prices and allocations that clear the date 0 market is fairly easy to compute. Each element (q_0, c_0) corresponds to an Arrow-Debreu equilibrium of an economy with two subjective expected utility maximizers. The whole set is then obtained by varying the parameter η over the unit interval. Recall that prices of consumption at nodes 4 and 5 are determined endogenously and may, in principle, vary across equilibria (in Figure 1 this is highlighted in the form of dashed branches). However, in this particular example, both agents assign unique probabilities to states ω_2 and ω_3 . As a consequence, consumption at the corresponding terminal nodes must have determinate relative prices. The set of market-clearing prices is therefore a ray:

$$q_0 \in \text{cone} \left\{ \left(1, \frac{1}{2}, \frac{1}{2}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4} \right) \right\}$$

At these prices the consumption plans of both agents as η varies are:

$$\begin{aligned} c_0^1 &= \left(\frac{9}{2}, \frac{1}{2} + \eta, \frac{3}{2} - \eta, 2\eta, 1, 1, 2 - 2\eta \right) \\ c_0^2 &= \left(\frac{9}{2}, 1, 1, 1, 1, 1, 1 \right) \end{aligned}$$

Agent 2 has complete preferences that result in a unique optimal plan at given prices q_0 . By contrast, Agent 1 has incomplete preferences that result in a continuum of undominated plans. These different plans do not involve different trading behavior in the market but are entirely implemented through different choices of production scales; in aggregate:

$$Y_0 = \left(-3, \frac{3}{2} + \eta, \frac{5}{2} - \eta, 1 + 2\eta, 0, 0, 3 - 2\eta \right)$$

Since the date 0 allocation c_0 serves as the status quo in date 1 markets, different choices of η must result in different equilibria. We pick as one particular instance $\hat{\eta} = \frac{1}{2}$, which results in the status quo $\hat{c}_0^1 = \hat{c}_0^2 = \left(\frac{9}{2}, 1, 1, 1, 1, 1, 1 \right)$. It is easy to verify that this allocation is Pareto efficient: The marginal rates cone of Agent 1 is

$$\nabla U^1[\hat{c}_0^1] = \text{cone} \left\{ \left(1, \frac{1+2\eta}{4}, \frac{3-2\eta}{4}, \frac{\eta}{2}, \frac{1}{4}, \frac{1}{4}, \frac{1-\eta}{2} \right) \mid \eta \in [0, 1] \right\},$$

and since the marginal rates cone of Agent 2 is the ray

$$\nabla U^2[\hat{c}_0^2] = \text{cone} \left\{ \left(1, \frac{1}{2}, \frac{1}{2}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4} \right) \right\},$$

it is true that $q_0 \in \nabla U^1[\hat{c}_0^1] \cap \nabla U^2[\hat{c}_0^2]$ and the necessary and sufficient condition (6) is met. 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 prior probabilities. If event ξ_1 is observed, updating leads to unique posterior probabilities




Figure 2: From left to right: Ambiguity partition (solid) before and after observing the event ξ_1 under maximum-likelihood updating.

$$\Pi_{\mathbb{M}}^1(\xi_1) = \left\{ \left(\frac{2}{3}, \frac{1}{3}, 0, 0 \right) \right\}, \quad \Pi_{\mathbb{M}}^2(\xi_1) = \left\{ \left(\frac{1}{2}, \frac{1}{2}, 0, 0 \right) \right\}.$$

For Agent 1, these are the conditional probabilities at $\eta = 1$, and thus the maximal probability mass is put on the observed event. By symmetry of the example, if ξ_2 is observed, conditional probabilities at $\eta = 0$ are taken, and the posterior probabilities are

$$\Pi_{\mathbb{M}}^1(\xi_2) = \left\{ \left(0, 0, \frac{1}{3}, \frac{2}{3} \right) \right\}, \quad \Pi_{\mathbb{M}}^2(\xi_2) = \left\{ \left(0, 0, \frac{1}{2}, \frac{1}{2} \right) \right\}.$$

Note that Knightian uncertainty is completely resolved at date 1. This is illustrated in Figure 2: The left panel shows the information partition \mathcal{I}_1 at date 1 (dashed), and the ambiguity partition \mathcal{A}^1 of the first agent. The large cell $\{\omega_1, \omega_4\}$ visualizes the ambiguity about the probability of these two states. The right panel shows the conditional ambiguity partition after observing event ξ_1 , which coincides with the information partition \mathcal{I}_2 at date 2. Just like the information partitions become finer over time, also the ambiguity partitions become finer over time – new information reduces uncertainty, both in the traditional and Knightian sense. However, this reduction in uncertainty has an effect on the decision problem of Agent 1. The first-order condition (5) implies that $q_0 \in \nabla U^1[\hat{c}_0^1]$. A section of this set inclusion along the $c_0^1(3)$ – $c_0^1(4)$ plane is depicted in the left panel of Figure 3. Knightian uncertainty results in a two-dimensional marginal rates cone. The right panel displays the same section after the information release: As ambiguity vanishes, the marginal rates cone collapses to a ray, and the price vector is no longer contained.

The consequences for the plan 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 \hat{c}_0^1 (more precisely, its section on the $c_0^1(3)$ – $c_0^1(4)$ 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 at node 3 and less consumption at node 4 than in the original plan \hat{c}_0^1 . If Agent 1 adjusts his demand, the market cannot clear under the original prices q_0 . Instead, market clearing results in different prices in the date 1 markets:




Figure 3: Section of Agent 1's budget set before and after the information release. The state price vector $(q_0(3), q_0(4))$ drops out of the marginal rates cone.

$$q_1 \in \text{cone} \left\{ \left(1, \frac{1}{2}, \frac{1}{2}, \frac{1}{4}, \frac{15}{76}, \frac{15}{76}, \frac{1}{4} \right) \right\}$$

Both agents react and choose different consumption plans under new prices q_1 :

$$\begin{aligned} \hat{c}_1^1 &= \left(\frac{9}{2}, \frac{18}{19}, \frac{18}{19}, \frac{24}{19}, \frac{4}{5}, \frac{4}{5}, \frac{24}{19} \right) \\ \hat{c}_1^2 &= \left(\frac{9}{2}, \frac{18}{19}, \frac{18}{19}, \frac{18}{19}, \frac{6}{5}, \frac{6}{5}, \frac{18}{19} \right) \end{aligned}$$

Agent 1 deviates from his initial plan because the resolution of Knightian uncertainty leads to dynamic inconsistency in his preferences. This dynamic inconsistency then spreads through the price mechanism and affects the choice set of Agent 2. As a consequence, the expected utility maximizer will also abandon his initial plan, in spite of having complete and dynamically consistent preferences. The pair (q_1, \hat{c}_1) clears the market at readjusted production levels

$$\hat{Y}_1 = \left(-3, \frac{36}{19}, \frac{36}{19}, \frac{42}{19}, 0, 0, \frac{42}{19} \right).$$

The new allocation is qualitatively different: To test it for Pareto efficiency, the necessary and sufficient condition (6) must be evaluated using the information available ex-ante, at date 0. The resulting marginal rates cone of Agent 1 is

$$\nabla U^1[\hat{c}_1^1] = \text{cone} \left\{ \left(1, \frac{38(1+2\eta)}{144}, \frac{38(3-2\eta)}{144}, \frac{57\eta}{144}, \frac{45}{144}, \frac{45}{144}, \frac{57(1-\eta)}{144} \right) \mid \eta \in [0, 1] \right\},$$

and its counterpart for Agent 2 is

$$\nabla U^2[\hat{c}_1^2] = \text{cone} \left\{ \left(1, \frac{76}{144}, \frac{76}{144}, \frac{38}{144}, \frac{30}{144}, \frac{30}{144}, \frac{38}{144} \right) \right\}.$$

A comparison of the marginal rates of substitution for consumption in nodes 4 and 5 reveals at one glance that $\nabla U^1[\hat{c}_1^1] \cap \nabla U^2[\hat{c}_1^2] = \emptyset$, and thus the allocation is Pareto inefficient. Since the market does not reopen at date 2, this is the final allocation, and the ensuing equilibrium is problematic from a normative viewpoint. Nevertheless, this equilibrium survives the equilibrium refinement: The tuple (\hat{q}, c) is an equilibrium with inertia. The source of inefficiency in this case is dynamic inconsistency of preferences.

Case 2. Full Bayesian updating: Now suppose both agents apply the full Bayesian rule. For Agent 2, this does not make a difference because the identities $\Pi_{\mathbb{B}}^2(\xi_1) = \Pi_{\mathbb{M}}^2(\xi_1)$ and $\Pi_{\mathbb{B}}^2(\xi_2) = \Pi_{\mathbb{M}}^2(\xi_2)$ hold trivially when priors are unique. However, for Agent 1, the set of posteriors upon observing ξ_1 now becomes the larger set

$$\Pi_{\mathbb{B}}^1(\xi_1) = \left\{ \left(\frac{2\eta_1}{3}, \frac{3-2\eta_1}{3}, 0, 0 \right) \mid \eta_1 \in [0, 1] \right\},$$

or its symmetric mirror image if ξ_2 is observed:

$$\Pi_{\mathbb{B}}^1(\xi_2) = \left\{ \left(0, 0, \frac{3-2\eta_2}{3}, \frac{2\eta_2}{3} \right) \mid \eta_2 \in [0, 1] \right\}.$$

Contrary to maximum-likelihood updating, there is no dynamic inconsistency. This is illustrated in Figure 4 for the case of observing ξ_1 . The marginal rates cones before and after the information release are displayed in the left and right panel respectively. These two cones agree. The full Bayesian rule updates probabilities in such a way that the first-order conditions at date 0 imply the first-order conditions at date 1, provided the agents stick to their initial plans.

Unfortunately, the full Bayesian rule leads to a different problem, which is illustrated in Figure 5. The initial ambiguity partition \mathcal{A}^1 visualizes the Knightian uncertainty between states ω_1 and ω_4 . There is no such ambiguity regarding the other states: Agent 1 is 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 $\mathcal{A}^1(\xi_1)$ is not finer than \mathcal{A}^1 ; instead, ambiguity has dilated and spread to the unambiguous event $\{\omega_2\}$. This is inconsistent with the principle that new information should not create new uncertainty – the agent forgets what he knew before. This affects the size of the equilibrium set: The two agents no longer agree on the probability of state ω_2 (or of state ω_3 if ξ_2 is observed). As a consequence, the relative price of consumption in this state is no longer unique, but the set of market clearing prices is enlarged. Dilation of ambiguity results in dilation of price indeterminacy. The status quo \hat{c}_0 gives rise to a continuum of equilibria with prices

$$q_1 \in \text{cone} \left\{ \left(1, \frac{1}{2}, \frac{1}{2}, \frac{1}{4}, \frac{27-12\eta_1}{60+16\eta_1}, \frac{27-12\eta_2}{60+16\eta_2}, \frac{1}{4} \right), \mid \eta_1, \eta_2 \in [0, 1] \right\}.$$




Figure 4: Section of Agent 1's budget set before and after the information release. The state price vector $(q_0(3), q_0(4))$ drops out of the marginal rates cone.




Figure 5: From left to right: Ambiguity partition (solid) before and after observing the event ξ_1 under full Bayesian updating.

and corresponding allocations after trading in the date 1 markets:

$$c_1^1 = \left(\frac{9}{2}, \frac{18}{15 + 4\eta_1}, \frac{18}{15 + 4\eta_2}, \frac{24\eta_1}{15 + 4\eta_1}, \frac{12 - 8\eta_1}{9 - 4\eta_1}, \frac{12 - 8\eta_2}{9 - 4\eta_2}, \frac{24\eta_2}{15 + 4\eta_2} \right)$$

$$c_1^2 = \left(\frac{9}{2}, \frac{18}{15 + 4\eta_1}, \frac{18}{15 + 4\eta_2}, \frac{18}{15 + 4\eta_1}, \frac{6}{9 - 4\eta_1}, \frac{6}{9 - 4\eta_2}, \frac{18}{15 + 4\eta_2} \right)$$

Again we pick one particular equilibrium from the continuum and test its final allocation for Pareto efficiency. The selection of $\hat{\eta}_1 = \hat{\eta}_2 = \frac{1}{2}$ results in consumption plans

$$\hat{c}_1^1 = \left(\frac{9}{2}, \frac{18}{17}, \frac{18}{17}, \frac{12}{17}, \frac{8}{7}, \frac{8}{7}, \frac{12}{17} \right)$$

$$\hat{c}_1^2 = \left(\frac{9}{2}, \frac{18}{17}, \frac{18}{17}, \frac{18}{17}, \frac{6}{7}, \frac{6}{7}, \frac{18}{17} \right).$$

This is again a deviation from the initial consumption plans, but this time the cause is not a change in preferences but a change in prices. This change in prices is consistent with

the definition of equilibrium because the resulting consumption plans satisfy the market clearing condition. The marginal rates cone of Agent 1 is

$$\nabla U^1[\hat{c}_1^1] = \text{cone} \left\{ \left(1, \frac{68(1+2\eta)}{288}, \frac{68(3-2\eta)}{288}, \frac{204\eta}{288}, \frac{63}{288}, \frac{63}{288}, \frac{204(1-\eta)}{288} \right) \mid \eta \in [0, 1] \right\},$$

and the marginal rates cone of Agent 2 is

$$\nabla U^2[\hat{c}_1^2] = \text{cone} \left\{ \left(1, \frac{136}{288}, \frac{136}{288}, \frac{68}{288}, \frac{84}{288}, \frac{84}{288}, \frac{68}{288} \right) \right\}.$$

This is the final allocation since there is no further trade at date 2, and again this allocation fails to be Pareto efficient: A simple inspection of marginal rates of substitution for nodes 4 and 5 reveals that $\nabla U^1[\hat{c}_1^1] \cap \nabla U^2[\hat{c}_1^2] = \emptyset$. This could not occur in the single-prior setting: If there is no ambiguous probability mass that could dilate, the no-trade theorem of Milgrom and Stokey (1982) applies, and agents would not deviate from a Pareto efficient status quo. The source of inefficiency in this case is dilation and the resulting expansion of price indeterminacy.

It should be noted that the equilibrium refinement is effective in this case: To go through as an equilibrium with inertia, the adjusted consumption plan \hat{c}_1^1 of Agent 1 would have to dominate the original plan \hat{c}_0^1 . To see that this is not the case, note that $\pi^1 = (\frac{1}{2}, \frac{1}{2}, 0, 0) \in \Pi_{\mathbb{B}}^1(\xi_1)$ and compare the resulting expected utility values. Since $\mathbb{E}^{\pi^1}[u^1(\hat{c}_0^1)] = 2 > 1.983 = \mathbb{E}^{\pi^1}[u^1(\hat{c}_1^1)]$, the adjusted plan does not dominate the status quo. In this case, inertia is a binding restriction. This suggests that the welfare effect of this equilibrium refinement depends on the chosen updating rules.

To conclude, both updating rules have side effects that are undesirable from a normative viewpoint. Full Bayesian updating may lead to dilation, and agents may fail to recall some of their previous knowledge. Maximum-likelihood updating may lead to dynamic inconsistency, and agents may no longer be satisfied with their original plans. Both effects have been shown to result in Pareto inefficient market outcomes, in spite of frictionless and complete markets. Moreover, these side effects may occur simultaneously: Dynamic inconsistency and dilation occur at the same time if the updating rule of Agent 1 is modified to $\Pi^1 = \alpha \Pi_{\mathbb{M}}^1 + (1 - \alpha) \Pi_{\mathbb{B}}^1$ for some weight $\alpha \in (0, 1)$.⁵ Do the presented 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 Information Processing

In this section we focus on normative considerations about the individual. Our aim is to identify updating rules that are consistent with individual rationality. In a setting with

⁵This is what Cheng (2021) calls *relative maximum-likelihood updating*.

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 affects behavior and knowledge over time. The *knowledge* of Agent i at date-event (t, ξ) is summarized by the tuple $(\xi, \Pi^i(\xi))$: The information set ξ contains all states that are considered possible, and thus represents uncertainty in the traditional sense. The set of posteriors $\Pi^i(\xi)$ contains all probability measures that are consistent with the observed event, and thus represents uncertainty in the Knightian sense. Section 3 introduced the concept of ambiguity partition for modeling the evolution of Knightian uncertainty over time. This kind of partition will prove useful for several of the results to come. Let pr_ζ denote the projection onto the subspace of those states contained in ζ .

Definition 4. The *ambiguity partition* $\mathcal{A}^i(\xi)$ of agent i , conditional on event ξ , is the finest partition of Ω that satisfies

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

The ambiguity partition is the finest partition that satisfies two properties: First, each cell is an *unambiguous* event; that is, any two measures $\pi^i, \hat{\pi}^i \in \Pi^i(\xi)$ assign identical probabilities $\pi^i|_{\mathcal{A}^i(\xi)} = \hat{\pi}^i|_{\mathcal{A}^i(\xi)}$ to all cells. This makes the power set $2^{\mathcal{A}^i(\xi)}$ a collection of unambiguous events in the sense of Epstein (1999). Second, the probability distribution within a cell is independent from the probability distribution within any other cell: For any two cells $\zeta, \zeta' \in \mathcal{A}^i(\xi)$, any two states $\omega \in \zeta$ and $\omega' \in \zeta'$, and any two measures $\pi^i, \pi^{i'} \in \Pi^i(\xi)$, there is some $\hat{\pi}^i \in \Pi^i(\xi)$ that assigns $\hat{\pi}^i(\omega) = \pi^i(\omega)$ and $\hat{\pi}^i(\omega') = \pi^{i'}(\omega')$. In the usual notation, $\mathcal{A}^i(\Omega) = \mathcal{A}^i$ is the unconditional ambiguity partition. A sequence of ambiguity partitions represents the evolution of uncertainty in the Knightian sense, just like the sequence of information partitions represents the evolution of uncertainty in the traditional sense. As a preliminary result, we shall establish that the concept of ambiguity partition is well-defined.

Proposition 1. *For any updating rule Π^i that satisfies Bayes' law, and any nonempty event $\xi \subseteq \Omega$, there exists a unique ambiguity partition conditional on ξ .*

Proof. Since Ω is a finite set, it admits a finite number of partitions. Let $\mathcal{P}(\Omega)$ denote the subset of partitions P that satisfy $\Pi^i(\xi) = \sum_{\zeta \in P} \text{pr}_\zeta(\Pi^i(\xi))$. By (1), the partition $P = \{\Omega\}$ is an element of $\mathcal{P}(\Omega)$, and therefore this set is nonempty. The relation *is finer than* defines a partial order on $\mathcal{P}(\Omega)$, and existence of a maximal element follows from the axiom of choice. Suppose this maximal element were not unique; then, there would be some state $\omega \in \Omega$, two ambiguity partitions $P \neq P'$ in $\mathcal{P}(\Omega)$, and two events $\zeta \in P$ and $\zeta' \in P'$ that satisfy $\zeta \neq \zeta'$ and the proper set inclusions $\{\omega\} \subset \zeta$ and $\{\omega\} \subset \zeta'$, for if any of these three conditions could not be met, one of the two partitions would either not be maximal or not distinct from the other. Since $\{\omega\}$ is not unambiguous, there must be two $\pi^i, \hat{\pi}^i \in \Pi^i(\xi)$ that satisfy $\pi^i(\omega) < \hat{\pi}^i(\omega)$. Since both ζ and ζ' are unambiguous, $\pi^i(\zeta) = \hat{\pi}^i(\zeta)$ and $\pi^i(\zeta') = \hat{\pi}^i(\zeta')$, but then the previous inequality implies both $\pi^i(\zeta \setminus \{\omega\}) > \hat{\pi}^i(\zeta \setminus \{\omega\})$ and $\pi^i(\zeta' \setminus \{\omega\}) > \hat{\pi}^i(\zeta' \setminus \{\omega\})$. Thus, the probabilities of

$\zeta \setminus \{\omega\}$ and $\zeta' \setminus \{\omega\}$ are not independent, and no proper subset of $\zeta \cup \zeta'$ could be a cell of an ambiguity partition. That is to say, P and P' could not have been ambiguity partitions in the first place – a contradiction. \square

Given that the set of possible updating rules is large, behavior and knowledge may respond to new information in many different ways. As the example in Section 3 shows, some of these responses are at odds with common sense. A normative desideratum is therefore to narrow down the set of rules on the ground of rationality. Clearly, normative arguments cannot be grounded on inductive reasoning based on examples. Behavior that seems surprising in a special case, need not be irrational in general. The correct approach is to define first the notion of rationality in the form of general criteria, in order to check then whether these are fulfilled by the updating rule under consideration. We focus on two such criteria.

4.1 Dynamic Consistency

The first rationality criterion is dynamic consistency: Since consumption in different states can be chosen independently, learning that a state has not occurred should not affect the ranking of consumption in other states; otherwise, preferences would vary with timing of information. This would lead to implausible judgment: Even though new information does not create new decision alternatives, previously undominated alternatives may become dominated. Dynamic consistency demands that the ranking $x \succ^i y$ of two state-contingent consumption plans is independent of how much information $\xi \subseteq \Omega$ about the state of the world the agents receives ex interim. Since updating is what connects conditional and unconditional preferences, this criterion of rationality is ultimately a requirement for updating rules.

Definition 5. An updating rule Π^i satisfies *dynamic consistency* if for any $\xi \in \mathcal{I}_t$, $t > 0$ and any x and y that agree on events disjoint from ξ the following condition is true:

$$x \succ^i y \iff x \succ_{\xi}^i y.$$

The example in Section 3 has demonstrated that prominent updating rules, such as maximum-likelihood updating, may fail to meet this requirement. Two alternatives that had been incomparable before could be ranked after an information release. Even though the choice set did not change, the agent found it optimal to deviate from the original plan. In principle, other updating rules may cause the same phenomenon, yet there is one exception: Bewley type preferences are dynamically consistent if and only if the full Bayesian rule (2) is used.

Proposition 2. *The unique updating rule that satisfies Bayes' law and dynamic consistency for any information structure \mathcal{I} is the full Bayesian rule:*

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




Figure 6: Problematic full Bayesian updating: The marginal distribution of the events (ζ_1, ζ_2) changes in spite of no news about their probabilities.

Proof. The proposition can be composed of previous results: That Bewley preferences are dynamically consistent under full Bayesian updating and Assumption 1 is the content of Lemma 8 of Kajii and Ui (2009). Uniqueness among all rules that satisfy Bayes' law is a by-product of the last sentence in Theorem 1 of Faro and Lefort (2019). \square

Proposition 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 dynamic consistency in the maxmin expected utility setting. By contrast, in the present setting such an updating rule exists, and it is identified in the proposition. Further restrictions, such as rectangularity, are not necessary.

4.2 Ambiguity Containmentment

The second rationality criterion is that agents do not forget what they knew. In the single-prior setting, this simply means that 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 multiple-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. A small example sheds light on the condition actually needed.

Example. Let $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ be the state space. Consider an agent with priors $\Pi^i = \{(\eta_1, \frac{1}{2} - \eta_1, \frac{1}{2} - \eta_2, \eta_2) \mid \eta_1, \eta_2 \in [0, \frac{1}{2}]\}$. The ambiguity partition $\mathcal{A}^i = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ is depicted in the left panel of Figure 6. Each cell is assigned a probability of $\frac{1}{2}$ under any prior in Π^i . Suppose the only payoff-relevant events are $\zeta_1 = \{\omega_1, \omega_2\}$ and $\zeta_2 = \{\omega_3, \omega_4\}$. Therefore, only the marginal distribution of (ζ_1, ζ_2) 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 $\xi = \{\omega_2, \omega_4\}$ 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 $\pi^i(\zeta_1|\xi) = \pi^i(\zeta_2|\xi) = 1/2$. However, if the agent applies the full Bayesian rule, he does not reach this conclusion. Updating Π^i prior by prior results in

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

and suddenly all choices are undominated. This dilation of ambiguity is illustrated in the right panel of Figure 6. It 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 $\pi^i = (0, 0, 0, 1)$, which puts all the mass on state ω_4 . 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 event ζ_1 , the agent finds it possible that the event occurs with a probability of zero. This is inconsistent with his priors, which assign an unambiguous probability of $1/2$. 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 two 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 $\Pi^i(\xi) = \{(0, 0, 0, 1)\}$. 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 contained 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 $\min_{\tilde{\pi}^i \in \Pi^i} \tilde{\pi}^i(\xi)$, 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 6. An updating rule satisfies *ambiguity containment* if for any $\xi \in \mathcal{I}_t$, $t > 0$ and any ambiguity partition \mathcal{A}^i 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 6 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 3. *Let $|\mathcal{A}| > 1$. The unique updating rule that satisfies Bayes' law and ambiguity containment for any information structure \mathcal{I} is the maximum-likelihood rule:*

$$\Pi^i(\xi) = \Pi_M^i(\xi) \quad \forall \xi \subseteq \Omega$$

Proposition 3, 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 the desired rationality criterion. Note that $|\mathcal{A}| > 1$ is not a limitation but a logical consistency condition: If the ambiguity partition contains only one cell, ambiguous probability mass cannot spill over to other cells by definition, and ambiguity containment holds trivially for any updating rule that satisfies Bayes' law.

4.3 A Difficulty in the Design of Updating Rules

Now that both criteria of individual rationality are formulated, it remains to be analyzed whether and when these criteria are fulfilled. 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; maximum-likelihood updating guarantees ambiguity containment. In light of these findings, a natural objective of normative theory is to design an updating rule that combines both criteria. Unfortunately, this is impossible:

Corollary 1. *Let $|\mathcal{A}| > 1$. There exists no updating rule that satisfies Bayes' law and guarantees both dynamic consistency and ambiguity containment for any information structure \mathcal{I} .*

Proof. Take Bayes' law as a primitive; then, the corollary follows trivially from Propositions 2 and 3 together with the observation from Section 3 that full Bayesian rule and maximum-likelihood rule result in different updates for some events. \square

Corollary 1 highlights a conflict between rationality 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, preferences 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, there are some events for which full Bayesian and maximum-likelihood updating agree. If an information structure only consists of such events, both rules coincide and thus exhibit dynamic consistency and ambiguity containment simultaneously. In particular, for arbitrary information structures, one can always find non-singleton sets of priors for which both rules lead to equivalent updates. The following two propositions provide sufficient and necessary conditions for unproblematic pairs of information structure and priors. These conditions depend 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 $\zeta \subseteq \xi$; that is, if the event fully contains the cell of the ambiguity partition. In this case, there is not even a difference between maximum-likelihood and minimum-likelihood updating:

Proposition 4. *If for each $t \in \{0, \dots, T - 1\}$,*

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

then for each $t \in \{0, \dots, T\}$,

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

Proof. Let $t \in \{0, \dots, T - 1\}$. Since \mathcal{A}^i is a partition, the probability of any event $\xi \in \mathcal{I}_t$ can be decomposed into

$$\pi^i(\xi) = \sum_{\zeta \in \mathcal{A}^i} \pi^i(\xi \cap \zeta).$$

For any two priors π^i and $\hat{\pi}^i$ in Π^i , $\xi \cap \zeta \in \{\zeta, \emptyset\}$ and Definition 4 jointly imply that $\pi^i(\xi \cap \zeta) = \hat{\pi}^i(\xi \cap \zeta)$ and thus $\pi^i(\xi) = \hat{\pi}^i(\xi)$. Therefore, $\Pi_{\mathbf{B}}^i(\xi) = \Pi_{\mathbf{M}}^i(\xi) = \Pi_{\mathbf{W}}^i(\xi) = \{\pi^i(\cdot | \xi)\}$. The same holds true at $t = T$ because \mathcal{I}_T consists of singletons, and for any updating rule, there is a unique element $\pi^i \in \Pi^i(\{\omega\})$, which satisfies $\pi^i(\omega) = 1$. \square

Proposition 4 covers as special cases the single-prior model, as \mathcal{A}^i consists of singletons in this case, and the two-date model, as $\{0, \dots, T - 1\} = \{0\}$ in this case. The next proposition complements this result with a necessary condition.

Proposition 5. *If for each $t \in \{0, \dots, T\}$,*

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

then for each $t \in \{0, \dots, T\}$,

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

Proof. Note that $\xi \cap \zeta = \zeta$ at date 0 because $\mathcal{I}_0 = \{\Omega\}$, and $\xi \cap \zeta = \xi$ at date T because \mathcal{I}_T contains only singletons. For all interim dates, the proof is by contradiction: If there is a nontrivial intersection of partitions, the updating rules must disagree on the observed event ξ involved. Lemma 1 in the appendix implies that for each $\zeta \in \mathcal{A}^i$ with $\xi \cap \zeta \neq \emptyset$, there is a unique scalar ν_ζ such that for any $\pi^{i*} \in \arg \max_{\pi^i \in \Pi^i} \pi^i(\xi)$, probabilities are $\pi^{i*}(\xi \cap \zeta) = \nu_\zeta$, and conditional probabilities are thus $\pi^{i*}(\zeta|\xi) = \frac{\nu_\zeta}{\pi^{i*}(\xi)}$. By (3), each element of $\Pi_{\mathbb{M}}(\xi)$ is of this form. Moreover, each scalar must satisfy $\nu_\zeta > 0$; otherwise, all priors would assign zero probability to states in $\xi \cap \zeta$, which violates Assumption 1. Suppose there were some ambiguity partition \mathcal{A}^i and some pair $(\xi, \zeta) \in \mathcal{I}_t \times \mathcal{A}^i$ for which $\xi \cap \zeta \notin \{\xi, \zeta, \emptyset\}$. Then, there must be some $\pi^i \in \Pi^i$ such that $\pi^i(\xi \cap \zeta) < \nu_\zeta$ but $\pi^i(\xi \cap \zeta') = \nu_{\zeta'}$ for any other $\zeta' \in \mathcal{A}^i \setminus \zeta$ with $\xi \cap \zeta' \neq \emptyset$. Note that $\pi^i(\xi) < \pi^{i*}(\xi)$ by construction. By (2), $\Pi_{\mathbb{B}}(\xi)$ contains the conditional probabilities $\pi^i(\cdot|\xi)$, which satisfy $\pi(\zeta'|\xi) = \frac{\nu_{\zeta'}}{\pi^i(\xi)} > \frac{\nu_{\zeta'}}{\pi^{i*}(\xi)}$ for each ζ' from before, but then these conditional probabilities cannot be contained in $\Pi_{\mathbb{M}}(\xi)$ – a contradiction. \square

Proposition 5 clarifies that differences between maximum-likelihood updating and full Bayesian updating only manifest if cells of information and ambiguity partitions have a nontrivial intersection. These insights generalize to a whole family of updating rules: Consider the family of *relative maximum-likelihood updating*, introduced by Cheng (2021). These rules are of the form $\Pi_\alpha^i(\xi) = \alpha \Pi_{\mathbb{M}}^i(\xi) + (1 - \alpha) \Pi_{\mathbb{B}}^i(\xi)$ for some scalar $\alpha \in [0, 1]$, and contain maximum-likelihood updating and full Bayesian updating as special cases. The following corollary is a useful by-product of the above analysis:

Corollary 2. *If for each $t \in \{0, \dots, T - 1\}$, $\xi \cap \zeta \in \{\zeta, \emptyset\} \quad \forall (\xi, \zeta) \in \mathcal{I}_t \times \mathcal{A}^i$, then each relative maximum-likelihood updating rule Π_α^i exhibits both dynamic consistency and ambiguity containment.*

Proof. Note that Π_α^i satisfies Bayes' law for any choice of $\alpha \in [0, 1]$. Proposition 4 implies that the whole continuum $\Pi_\alpha^i(\xi) = \alpha \Pi_{\mathbb{M}}^i(\xi) + (1 - \alpha) \Pi_{\mathbb{B}}^i(\xi)$ of updating rules collapses to a singleton when restricted to cells of \mathcal{I}_t , $t \geq 0$. The corollary now follows from Propositions 2 and 3. \square

Corollary 2 makes relative maximum-likelihood updating testable, at least within the framework of Knightian decision theory. If ambiguity and information structure are designed to meet the above intersection criterion, then the entire family of updating rules can be rejected in the laboratory if dynamic inconsistency or dilation is observed. Both properties have been tested for in experiments, and indeed subjects exhibit behavior that

can be rationalized by maximum-likelihood or full Bayesian updating: Testing for dynamic consistency, Dominiak, Duersch, and Lefort (2012) find that about 68% of subjects make dynamically inconsistent decisions in a version of the Ellsberg urn experiment. Shishkin and Ortoleva (2021) test for dilation, albeit under the assumption of maxmin expected utility, but find no evidence, which is at best consistent with maximum-likelihood updating. Also assuming maxmin expected utility, Ngangoué (2021) finds more diverse updating in a simulated asset market experiment with an information structure identical to the one from the example in Section 3. She identifies the behavior of 54% of subjects as consistent with full Bayesian updating and of further 35% as consistent with maximum-likelihood updating. These numbers are in the same ball park as findings of Cohen, Gilboa, Jaffray, and Schmeidler (2000) in a dynamic Ellsberg urn experiment. The authors classify about 55% of subjects as full Bayesian updaters, and further 29% as maximum-likelihood updaters.

4.4 Rectangularity

The true value of our intersection criterion becomes manifest once it is related 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 7. A tuple (\mathcal{I}, Π^i) of information structure and subjective probabilities satisfies *rectangularity* if for each $t \in \{1, \dots, T-1\}$ and each $\xi \in \mathcal{I}_{t-1}$,

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

Since Definition 7 is based on the full Bayesian rule, $\Pi_{\mathbf{B}}^i$, 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 full Bayesian, maximum-likelihood, and minimum-likelihood updating disappear. In fact, the concurrence of these three updating rules is the very nature of rectangularity:

Theorem 1. *The following statements are equivalent:*

(i) *For each $t \in \{0, \dots, T\}$,*

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

(ii) *(\mathcal{I}, Π^i) satisfies rectangularity.*

Theorem 1, whose proof can be found in the appendix, reveals that the relevance of rectangularity is not limited to maxmin expected utility. In the Knightian decision model, it guarantees dynamic consistency jointly with ambiguity containment whenever

the agent uses one of its defining updating rules. Moreover, in combination with Proposition 5, the theorem suggests an indirect test for rectangularity that is not confined to a particular decision model. Rectangularity can be rejected purely on the basis of the ambiguity partition, which is easier to elicit in the laboratory than the entire set of subjective probabilities. If the ambiguity partition exhibits nontrivial intersection with an information partition, subjective probabilities can by no means be rectangular. Setting up a dynamic decision problem is not necessary for this simple test.

In light of Theorem 1, it becomes clear that Definition 7 can be equivalently formulated in terms of maximum-likelihood updates by replacing $\Pi_{\mathbf{B}}^i$ with $\Pi_{\mathbf{M}}^i$. The resulting condition can be strengthened as follows:

Definition 8. A tuple (\mathcal{I}, Π^i) of information structure and subjective probabilities satisfies *maximum rectangularity* if for each $t \in \{1, \dots, T-1\}$ and each $\xi \in \mathcal{I}_{t-1}$,

$$\Pi_{\mathbf{M}}^i(\xi) = \bigcup_{\pi^i \in \Delta^{|\Omega^i|}} \sum_{\xi' \in \mathcal{I}_t} \pi^i(\xi'|\xi) \Pi_{\mathbf{M}}^i(\xi').$$

The main difference to Definition 7 is that the union is not taken over Π^i but over the unit simplex $\Delta^{|\Omega^i|}$, which is a larger set. This is where the formulation in terms of the maximum-likelihood updating matters: Under maximum rectangularity the set Π^i contains boundary points of the simplex that assign a probability of zero to some observable events. This is at odds with the second sentence of Assumption 1 but unproblematic when the maximum-likelihood rule is used. Since this rule only updates priors that put a maximal of probability mass on the observed event, the denominator in (1) is always nonzero. By contrast, under full Bayesian and minimum-likelihood updating, some conditional probabilities would be undefined in this case. For this reason, Definition 8 goes hand in hand with a relaxation of Assumption 1.

5 Welfare

In this section we focus on normative considerations about society. Our aim is to identify equilibria that engender social welfare. As we use Pareto efficiency at the ex-ante stage as our welfare standard, the question we are implicitly asking is whether some allocation mechanism other than a sequence of competitive markets could make every agent better off. Since no restrictions are placed on the set of alternative mechanisms, this set includes direct allocation by a planner with perfect knowledge about preferences and endowments. Recall that complete contingent markets result in efficient allocations in the two-date setting: Rigotti and Shannon (2005) show that date 0 markets always have a Pareto efficient outcome, and if there is no more trade at date 1, this efficient outcome is the final equilibrium allocation. The same logic can be extended to the setting with multiple future dates: An equilibrium is Pareto efficient if agents never deviate from the allocation attained in date 0 markets. This is where the first rationality criterion from Section 4

comes into play: Dynamic consistency and inertia jointly imply that agents stick to their initial plans. This insight leads to the following theorem.

Theorem 2. *If (\mathcal{I}, Π^i) satisfies dynamic consistency for each agent i , then all equilibria with inertia are Pareto efficient.*

Proof. At any equilibrium with inertia (q, c) , the date 0 allocation $c_0 = (c_0^1, \dots, c_0^I)$ satisfies the first-order conditions (5) for undominatedness. Since this implies the characterization (6) of Pareto efficiency, c_0 is an ex-ante efficient allocation under the unconditional preferences $(\succ^1, \dots, \succ^I)$. The remainder of the proof is by induction: We start at $t = 1$, and show that if c_{t-1} is Pareto efficient, then c_t must be Pareto efficient as well. Irrespective of prices q_t at the subsequent date, each status quo c_{t-1}^i is contained in the budget set $B_t(q_t, c_{t-1}^i)$. By Definition 3, $c_t \neq c_{t-1}$ only if c_t Pareto dominates c_{t-1} under the interim preferences $(\succ_t^1, \dots, \succ_t^I)$. By Proposition 2, dynamic consistency implies that all updating rules agree with full Bayesian updating on \mathcal{I} . By Proposition 4 of Kajii and Ui (2009), full Bayesian updating implies that the ex-ante efficient allocation c_{t-1} is also interim efficient, and thus $c_t = c_{t-1}$ by inertia. Since this is true for any $t > 0$, it follows that $c_0 = \dots = c_{T-1}$, and thus the terminal allocation c_{T-1} is Pareto efficient. \square

In the proof of Theorem 2, we take a huge shortcut by utilizing the characterization of interim efficient allocations of Kajii and Ui (2009). Since Bewley's inertia assumption implies that agents only deviate from the status quo if the new allocation is conditionally preferred, this will never happen unless the status quo is interim Pareto dominated. If updates are consistent with the full Bayesian rule, such interim Pareto domination does not occur, and all agents stick to the original plan. This line of reasoning is simple because prices $q = (q_0, \dots, q_{T-1})$ can be disregarded. For this reason, it should be emphasized that Theorem 2 does not imply that prices stay constant over time. The part of equilibrium (q, c) that inertia keeps constant is the sequence of consumption plans $c_0 = c_1 = \dots = c_{T-1}$. Nevertheless, bounds are placed on the variance of prices over time by the size of each (conditional) marginal rates cone and the individual production technologies. This follows from the first-order condition (5).

Finally, we will drop the inertia assumption. One of the main insights from Section 3 is that prices and consumption plans can change in such a way that agents abandon an efficient allocation. This may happen in spite of dynamic consistency because budget sets change over time. The idea behind the next theorem is not preventing such changes of prices and plans but ensuring that plans only alternate within the set of efficient allocations. A necessary condition suggested by the example in Section 3 is that ambiguous probability mass does not dilate. This insight is strengthened by the following theorem: If priors are maximum-rectangular, ambiguity containment is not only necessary but also sufficient.

Theorem 3. *If (\mathcal{I}, Π^i) satisfies ambiguity containment and maximum rectangularity for each agent i , then all equilibria are Pareto efficient.*

Theorem 3 utilizes the fact that the Pareto frontier is a large set when preferences are incomplete. Its proof, which can be found in the appendix, is based on the relation between the first-order conditions under conditional preferences and the characterization of Pareto efficiency. Maximum rectangularity of subjective probabilities translates into maximal rectangularity of marginal rates cones. If a price vector q_t is contained in the conditional marginal rates cones for all observable events $\xi \in \mathcal{I}_t$ at that date, then q_t must also be contained in the unconditional marginal rates cone of the agent. It should be noted that additive separability of utility is a relevant property: Since past prices cannot change at a future date, marginal utility of past consumption must not react to changes in future consumption either; otherwise, marginal rates and relative prices of past consumption would cease to agree.

In light of Theorem 1, it should not be surprising that preferences are dynamically consistent under the assumptions of Theorem 3. However, with inertia dropped, there is no reason to believe that agents stick to their original consumption plans. This is because preferences are incomplete. Deviations from the status quo to an incomparable plan are always possible, even in cases where prices $q_0 = q_1 = \dots = q_{T-1}$ are constant across time. The theorem shows that repeated trade is not inconsistent with Pareto efficiency. In fact, a propensity to trade can be viewed as a natural feature of Knightian uncertainty.

6 Conclusion

The traditional theory of general equilibrium holds strong implications for complete financial markets: Individuals plan their lifetime consumption in advance, trade only once, and stick to their plans when markets reopen. While appealing from a welfare-theoretic viewpoint – the resulting allocations are Pareto efficient – these predictions are contradicted by the empirical fact that market participants trade repeatedly. One possible explanation for this behavior is Knightian uncertainty. With imprecise knowledge of probabilities, the dynamics of individual behavior and market outcomes depend on the updating rule. This is where decision theory and general equilibrium theory meet. Normative criteria for the individual imply normative criteria of social welfare.

As regards the individual, we expect that preferences are dynamically consistent, and that better information reduces uncertainty. These two rationality criteria are not always compatible. Even though we have identified updating 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. In some economies, one of the two properties must be given up, but neither is a natural candidate. Dynamic consistency is a form of commitment to one's own plans. In absence of such intrinsic commitment, individuals might be better off if their market participation were restricted sometimes, which does not suit with the assumption of market completeness.

Allowing uncertainty to spread, on the other hand, 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: Individuals should use information to improve their knowledge, and should not forget knowledge previously held.

As regards society, our normative criterion is Pareto efficiency. This criterion turns out to be demanding: Even in complete markets, it cannot be guaranteed unless the both rationality criteria are met. If either of the two 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 rationality criteria 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 subjective probabilities are updated, which is ultimately an empirical question.

Appendix

The appendix contains several proofs that have been omitted in the main part of this paper. It will be notationally convenient to decompose updating rules into projections and rescale operations. For any set $S \subset \mathbb{R}^{|\Omega|}$, define the rescale operation $R(S) = \left\{ \frac{x}{x \cdot \vec{1}} \mid x \in S \right\}$, the maximizer $M_\xi(S) = \arg \max_{x \in S} x \cdot \text{pr}_\xi(\vec{1})$, and the minimizer $W_\xi(S) = \arg \min_{x \in S} x \cdot \text{pr}_\xi(\vec{1})$. The three updating rules can be written as

$$\Pi_{\mathbb{B}}^i(\xi) = R(\text{pr}_\xi(\Pi^i)) \tag{7}$$

$$\Pi_{\mathbb{M}}^i(\xi) = R(\text{pr}_\xi(M_\xi(\Pi^i))) \tag{8}$$

$$\Pi_{\mathbb{W}}^i(\xi) = R(\text{pr}_\xi(W_\xi(\Pi^i))). \tag{9}$$

We first introduce two lemmata that will play a role in the proofs to come.

Lemma 1. *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

Lemma 2. *Let (q, c) be an equilibrium. For each $t \geq 0$, there is some $Y_t \in \bar{\mathcal{Y}} \cap \mathcal{C}_t(Y_{t-1})$ that satisfies $q_t \in N_{\bar{\mathcal{Y}}}[Y_t]$ (with $\mathcal{C}_0(Y_{-1}) = \mathcal{C}$).*

Proof of Proposition 3. Necessity of Π_M^i : By Lemma 1, 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 $\kappa \neq 0$ such that

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

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,

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

and we have $\kappa = 1/\pi^{i*}(\xi)$. 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 1 and establishes ambiguity containment. \square

Proof of Theorem 1. Both implications are proven separately. The first part of this proof establishes that (i) implies (ii): As a prelude, note that the following inclusion always holds, irrespective of (i), by Bayes' law:

$$\Pi_{\mathbb{B}}^i(\xi) \supseteq \bigcup_{\pi^i \in \Pi^i} \sum_{\xi' \in \mathcal{I}_t} \pi^i(\xi'|\xi) \Pi_{\mathbb{B}}^i(\xi'). \quad (10)$$

To see this, note that for any $\pi^i \in \Pi_{\mathbb{B}}^i(\xi)$, the definition of conditional probabilities (1) implies that $\pi^i = \sum_{\xi' \in \mathcal{I}_t} \pi^i(\xi') \pi^i(\cdot|\xi')$. Therefore, $\pi^i \in \sum_{\xi' \in \mathcal{I}_t} \pi^i(\xi') \Pi_{\mathbb{B}}^i(\xi')$, and (10) follows since $\pi^i(\xi') = \pi^i(\xi'|\xi)$ because $\pi^i(\xi) = 1$. We now use (i) to establish the opposite inclusion

$$\Pi_{\mathbb{B}}^i(\xi) \subseteq \bigcup_{\pi^i \in \Pi^i} \sum_{\xi' \in \mathcal{I}_t} \pi^i(\xi'|\xi) \Pi_{\mathbb{B}}^i(\xi') \quad (11)$$

by contradiction. Suppose some $\pi^i \notin \Pi_{\mathbb{B}}^i(\xi)$ were a member of the set on the right-hand side of (11) for some t . Pick an arbitrary $\xi' \in \mathcal{I}_t$ and consider the projection $p^i = \text{pr}_{\xi'}(\pi^i)$; then, $R(p^i) \in \Pi_{\mathbb{B}}(\xi')$. We use the fact that full Bayesian updating is transitive for $\xi' \subseteq \xi$:

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

Together with (i) this yields $\Pi_{\mathbb{M}}^i(\xi') = \Pi_{\mathbb{W}}^i(\xi') = R(\text{pr}_{\xi'}(\Pi_{\mathbb{B}}^i(\xi)))$. By construction, $p^i \in \text{pr}_{\xi'}(\Pi_{\mathbb{B}}^i(\xi))$, and thus there must be two other points $\bar{p}^i \in \text{pr}_{\xi'}(M_{\xi'}(\Pi_{\mathbb{B}}^i(\xi)))$ and $\underline{p}^i \in \text{pr}_{\xi'}(W_{\xi'}(\Pi_{\mathbb{B}}^i(\xi)))$ that satisfy $R(p^i) = R(\bar{p}^i) = R(\underline{p}^i)$. Since the preimage of a point under R is a ray, there must be some $\alpha \in [0, 1]$ such that $p^i = \alpha \bar{p}^i + (1 - \alpha) \underline{p}^i$. As the domain restriction of $\text{pr}_{\xi'}$ to the unit simplex $\Delta^{|\Omega|}$ is injective, there are unique elements $\bar{\pi}^i = \text{pr}_{\xi'}^{-1}(\bar{p}^i) \in M_{\xi'}(\Pi_{\mathbb{B}}^i(\xi))$ and $\underline{\pi}^i = \text{pr}_{\xi'}^{-1}(\underline{p}^i) \in W_{\xi'}(\Pi_{\mathbb{B}}^i(\xi))$. Both $\bar{\pi}^i$ and $\underline{\pi}^i$ are elements of $\Pi_{\mathbb{B}}^i(\xi)$. Since this set is convex under Assumption 1, the convex combination $\hat{\pi}^i = \alpha \bar{\pi}^i + (1 - \alpha) \underline{\pi}^i$ is also an element of $\Pi_{\mathbb{B}}^i(\xi)$. By linearity of the projection $\text{pr}_{\xi'}$, $\hat{\pi}^i = \text{pr}_{\xi'}^{-1}(\alpha \bar{p}^i + (1 - \alpha) \underline{p}^i) = p^i$, but this means that $\pi^i \in \Pi_{\mathbb{B}}^i(\xi)$ – a contradiction. As this argument is independent of the particular choice of $t > 0$ and $\xi' \in \mathcal{I}_t$, the inclusion (11) is established, and the first part of this proof is complete.

Now we move on to the second part of this proof and show that (ii) implies (i): We show that for any $t > 0$, a recursive application of the rectangularity condition (Definition 7) from date 0 through date $t - 1$ implies that $\Pi_{\mathbb{B}}^i$ and $\Pi_{\mathbb{M}}^i$ agree on all date t events (the arguments for $\Pi_{\mathbb{W}}^i$ will be analogous): 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

$$\Pi^i = \bigcup_{\pi_1^i \in \Pi^i} \sum_{\xi_1 \in \mathcal{I}_1} \pi_1^i(\xi_1) \bigcup_{\pi_2^i \in \Pi^i} \sum_{\xi_2 \in \mathcal{I}_2} \pi_2^i(\xi_2 | \xi_1) \cdots \bigcup_{\pi_t^i \in \Pi^i} \sum_{\xi_t \in \mathcal{I}_t} \pi_t^i(\xi_t | \xi_{t-1}) \Pi_{\mathbb{B}}^i(\xi_t). \quad (12)$$

For any $\xi \in \mathcal{I}_t$, the maximum operation M_ξ can be applied to both sides of (12). 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 (12) 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 (12) thus leads to

$$M_\xi(\Pi^i) = \bigcup_{\pi_1^i \in M_{\xi_1}(\Pi^i)} \sum_{\xi_1 \in \mathcal{I}_1} \pi_1^i(\xi_1) \cdots \bigcup_{\pi_t^i \in M_{\xi_t}(\Pi^i)} \sum_{\xi_t \in \mathcal{I}_t} \pi_t^i(\xi_t | \xi_{t-1}) \Pi_{\mathbb{B}}^i(\xi_t). \quad (13)$$

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 (13) results in

$$\text{pr}_\xi(M_\xi(\Pi^i)) = \bigcup_{\pi_1^i \in M_{\xi_1}(\Pi^i)} \cdots \bigcup_{\pi_t^i \in M_{\xi_t}(\Pi^i)} \text{pr}_\xi \left(\pi_1^i(\hat{\xi}_1) \pi_2^i(\hat{\xi}_2 | \hat{\xi}_1) \cdots \pi_t^i(\hat{\xi}_t | \hat{\xi}_{t-1}) \Pi_{\mathbb{B}}^i(\xi) \right), \quad (14)$$

and the rescale operation R applied to both sides of (14) in

$$R(\text{pr}_\xi(M_\xi(\Pi^i))) = R(\lambda^* \text{pr}_\xi(\Pi_{\mathbb{B}}^i(\xi))). \quad (15)$$

By construction, $R(\lambda S) = R(S)$ for any scalar $\lambda \neq 0$ and set S . Therefore, by means of (7) and (8), Equation (15) can be rewritten as

$$\Pi_{\mathbb{M}}^i(\xi) = \Pi_{\mathbb{B}}^i(\xi). \quad (16)$$

Since this is true for any $t > 0$ and any choice of $\xi \in \mathcal{I}_t$, the agreement of $\Pi_{\mathbb{B}}^i$ and $\Pi_{\mathbb{M}}^i$ is proven. The agreement of $\Pi_{\mathbb{B}}^i$ and $\Pi_{\mathbb{W}}^i$ is obtained by the same arguments after replacing the maximum operation M_ξ with the minimum operation W_ξ . This establishes the second implication and concludes the proof. \square

Proof of Theorem 3. By Lemma 2, $q_t \in N_{\bar{\mathcal{Y}}}[Y]$ for some $Y \in \bar{\mathcal{Y}}$ at any date $t \geq 0$. If $q_t \in \nabla U^i[c_t^i]$ for any date $t \geq 0$ and consumer i , then the characterization (6) of Pareto

efficiency is satisfied throughout, and the market outcome at any date is efficient. To establish this property in compact notation, define the binary operator $\otimes : \mathbb{R}^L \times \mathbb{R}^{|\Omega|} \rightarrow \mathbb{R}^L$ as $q_t \otimes \pi^i = \left((q_t(t, \xi) \pi^i(\xi))_{\xi \in \mathcal{I}_t} \right)_{t=0}^T$. As u^i is differentiable under Assumption 2, this operator can be used to decompose the marginal rates cone $\nabla U^i[c_0^i] = \text{cone}(v_0^i \otimes \Pi^i)$, into a marginal utility vector $v_0^i \in \mathcal{C}_{++}$ and the set of beliefs. Taking Lemma 2 into account, the first-order condition (5) at date zero 0 boils down to

$$q_0 \in \text{cone}(v_0^i \otimes \Pi^i).$$

At any later date $t > 0$, $v_t^i = \mathcal{C}_t(v_{t-1}^i)$ because $c_t^i = \mathcal{C}_t(c_{t-1}^i)$ and marginal utility of past consumption is invariant to changes in future consumption since u^i is additively separable under Assumption 2. Defining the projection $\text{pr}_\xi : \mathcal{C} \rightarrow \mathcal{C}$ in a slight reuse of notation as $\text{pr}_\xi(q_t) = \left((q_t(t, \xi') \mathbb{1}_{\{\xi' \subseteq \xi\}}(\xi'))_{\xi' \in \mathcal{I}_t} \right)_{t=0}^T$, the first-order conditions for undominatedness under the conditional preference relation \succ_ξ^i can be written as

$$\text{pr}_\xi(q_t) \in \text{cone}(v_t^i \otimes \Pi_M^i(\xi)),$$

in which posterior beliefs are computed by means of the maximum-likelihood rule because this is necessary for ambiguity containment according to Proposition 3. These conditions can be aggregated across all $\xi \in \mathcal{I}_t$ into

$$\begin{aligned} q_t &\in \sum_{\xi \in \mathcal{I}_t} \text{cone}(v_t^i \otimes \Pi_M^i(\xi)) = \bigcup_{\alpha \in \mathbb{R}_+^{|\mathcal{I}_t|}} \sum_{\xi \in \mathcal{I}_t} \alpha_\xi (v_t^i \otimes \Pi_M^i(\xi)) \\ &= v_t^i \otimes \bigcup_{\alpha \in \mathbb{R}_+^{|\mathcal{I}_t|}} \sum_{\xi \in \mathcal{I}_t} \alpha_\xi \Pi_M^i(\xi) = \text{cone} \left(v_t^i \otimes \bigcup_{\pi^i \in \Delta^{|\Omega|}} \sum_{\xi \in \mathcal{I}_t} \pi^i(\xi) \Pi_M^i(\xi) \right), \end{aligned}$$

in which v_t^i could be shifted outside the summation because \otimes is a linear operator. Since maximum rectangularity is satisfied and since $\Pi_M^i(\Omega) = \Pi^i$, Definition 8 can be used to rewrite the term inside the conical hull in order to obtain

$$q_t \in \text{cone}(v_t^i \otimes \Pi^i),$$

and thus $q_t \in \nabla U^i[c_t^i]$. Since this holds for any date $t \geq 0$ and any consumer i , (6) is satisfied and c_0, \dots, c_{T-1} is a sequence of Pareto efficient allocations. \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. 1987. Knightian Decision Theory. Part II: Intertemporal Problems. Discussion Paper No. 835, Cowles Foundation.
- Bewley, T. F. 2002. Knightian Decision Theory. Part I. *Decisions in Economics and Finance* 25 (2), 79–110.
- Chambers, R. G. 2014. Uncertain Equilibria and Incomplete Preferences. *Journal of Mathematical Economics* 55, 48–54.
- Cheng, X. 2021. Relative Maximum Likelihood Updating of Ambiguous Beliefs. *Journal of Mathematical Economics* forthcoming.
- 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.
- Faro, J. H., and Lefort, J.-P. 2019. Dynamic Objective and Subjective Rationality. *Theoretical Economics* 14 (1), 1–14.
- Gerasimou, G. 2018. On the Indifference Relation in Bewley Preferences. *Economics Letters* 164, 24–26.
- 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.
- Kajii, A., and Ui, T. 2009. Interim Efficient Allocations under Uncertainty. *Journal of Economic Theory* 144 (1), 337–353.
- Knight, F. H. 1921. *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.
- Milgrom, P., and Stokey, N. 1982. Information, Trade and Common Knowledge. *Journal of Economic Theory* 26 (1), 17–27.

- Ngangoué, K. 2021. Learning under Ambiguity: An Experiment in Gradual Information Processing. *Journal of Economic Theory* 195, 1–30.
- Rigotti, L., and Shannon, C. 2005. Uncertainty and Risk in Financial Markets. *Econometrica* 73 (1), 203–243.
- Savage, L. J. 1954. *The Foundations of Statistics*. New York: Wiley.
- Seidenfeld, T., and Wasserman, L. 1993. Dilation for Sets of Probabilities. *Annals of Statistics* 21 (3), 1139–1154.
- Shishkin, D., and Ortoleva, P. 2021. Ambiguous Information and Dilation: An Experiment. working paper.