

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays in Mathematical Economics

Permalink

<https://escholarship.org/uc/item/8673j381>

Author

Vasquez, Markus Antonio

Publication Date

2017

Peer reviewed|Thesis/dissertation

Essays in Mathematical Economics

by

Markus Antonio Vasquez

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Robert M. Anderson, Chair

Professor David S. Ahn

Professor Christina M. Shannon

Spring 2017

Essays in Mathematical Economics

Copyright 2017
by
Markus Antonio Vasquez

Abstract

Essays in Mathematical Economics

by

Markus Antonio Vasquez

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Robert M. Anderson, Chair

We apply mathematical techniques in the context of economic decision making. First, we are interested in understanding the behaviors and beliefs of agents playing economic games in which the underlying action spaces are possibly non-compact and the agents' payoff functions are possibly discontinuous. Under these circumstances, there is no guarantee of the existence of a Nash equilibrium in randomized strategies. In fact, there are games for which no Nash equilibrium exists. To restore equilibrium we allow each agent access to randomized strategies that are not necessarily countably additive. This has the unfortunate side effect of introducing uncertainty into the players' payoff functions due to the failure of Fubini's theorem for finitely additive measures. We introduce two ways of resolving this ambiguity and show that for one we are able to recover a general equilibrium existence result.

Next, we turn to the problem that expected utility theory typically assumes that agents use concave utility functions. This is problematic since this implies that agents are risk averse and, consequently, will not gamble. We speculate that non-concavity may be the result of agents' utility functions arising from solving the knapsack problem, a combinatorial optimization problem. We introduce a class of utility of wealth functions, called knapsack utility functions, which are appropriate for agents who must choose an optimal collection of indivisible goods from a countably infinite collection. We find that these functions are pure jump processes. Moreover, we find that localized regions of convexity—and thus a demand for gambling—is the norm, but that the incentive to gamble is much more pronounced at low wealth levels. We consider an intertemporal version of the problem in which the agent faces a credit constraint. We find that the agent's utility of wealth function closely resembles a knapsack utility function when the agent's saving rate is low.

Finally, we turn our attention to the Black-Scholes model of security price movements. Our goal is to understand the beliefs and incentives of individual agents required for the Black-Scholes model to be self-predicting. We consider a model in which each agent believes that the Black-Scholes model is correct. Each agent observes a private stream of information, which she uses to update her beliefs about future movements of the security price. Each agent is then faced with an optimization problem whose solution tells us her optimal portfolio for any given price (i.e. her demand function). Imposing market clearing conditions then

determines a price at each point in time. That is, the agents prior beliefs about the security price process along with their private information streams generate a price process. We may then ask under which conditions the distribution of this process matches the agents' prior belief. We find that that condition is fairly restrictive and imposes significant constraints on the drift of the price process when agents are homogenous and use utility functions with constant absolute risk aversion or constant relative risk aversion.

To Mark, Laurie and Angela Vasquez

Contents

Contents	ii
List of Figures	iv
1 Equilibrium with finitely additive mixed strategies	1
1.1 Introduction	1
1.2 The Bertrand-Edgeworth model of duopoly pricing with limited supply . . .	2
1.3 A game with no value	4
1.4 The Bertrand competition model of duopoly pricing with anti-predatory pricing regulation	6
1.5 Mathematical preliminaries	7
1.6 Game theoretic setup	13
1.7 Optimistic equilibria	17
1.8 Vague equilibria	19
1.9 Relations between equilibrium concepts	21
1.10 Agreeable functions	23
1.11 Conclusion	26
2 Utility of wealth with many indivisibilities	27
2.1 Introduction	27
2.2 Knapsack utility functions	28
2.3 Insurance and gambling	30
2.4 Measures of risk seeking	32
2.5 Approximate concavity at large wealth levels	36
2.6 Proof that knapsack utility functions are pure jump processes	38
2.7 An intertemporal model with credit constraints	41
2.8 Conclusion	44
3 Information Aggregation in Financial Markets	46
3.1 Introduction	46
3.2 The model	47
3.3 CARA utility functions	48
3.4 CRRA utility functions	52

3.5 Conclusion	56
Bibliography	58

List of Figures

1.1	A game with no value	5
2.1	A knapsack utility function and its convex hull.	30
2.2	The agent begins with some non-random allocation, point A. Exposure to a contingent liability, moves her to point B. Purchasing an insurance policy with a positive premium moves her to point C. Finally, accepting a negative mean gamble moves her to point D. The agent would have obtained the same increase in expected utility by moving directly from B to D by simultaneously gambling and buying insurance.	33
2.3	At wealth levels 5, 15, 25, . . . the agent will be able to achieve $\bar{U}(5)$, $\bar{U}(15)$, $\bar{U}(25)$, . . . by betting 5 on a fair coin flip. Accepting the gamble increases the agent's expected utility at these wealth levels, but the amount by which the expected utility increases tends to 0 as the wealth level becomes large.	37
2.4	If the agent values utility a year from now at about 90% of current utility, we see that V_s very close to U when s is small is approximately concave when s is large.	43
3.1	Sample paths of r with CARA utility functions, $r_0 = 0.05$, and $\sigma = 0.1$	49
3.2	Sample paths of r with CRRA utility functions, $r_0 = 0.05$, and $\sigma = 0.1$	53

Acknowledgments

First and foremost, I would like to thank my advisor Bob Anderson for years of weekly encouragement and guidance. I don't imagine that I was an easy mentee, but Bob's patient suggestions helped me focus my energies on fruitful projects even when I was emotionally invested in something that simply wasn't panning out. During periods of discouragement, Bob was able to strike the right balance of gentle prodding and fatherly reassurance to help me move on. My choice of advisor was one the best choices that I made during my time in graduate school and seemed to compensate for any missteps.

I would like to thank Chris Shannon for serving on my dissertation committee and for providing extensive feedback on my work on knapsack utility functions.

I would like to thank David Ahn for taking the time to serve first on my qualifying exam committee and then again on my dissertation committee.

I would like to thank Hugh Woodin and David Aldous for serving on my qualifying exam committee.

I would like to thank George Mailath, Simon Grant, and several anonymous referees for helpful suggestions and for reading drafts of a paper related to the second chapter of this dissertation.

I would like to thank Barb Waller and Vicky Lee for helping me navigate the various complex systems that make up the university.

I would like to thank my parents, Mark and Laurie for always being there to catch me when I fell and for always believing in me.

I would like to thank my mother Angela for teaching me to read, to learn, to stand on my own two feet and to play outside.

I would like to thank my Aunt Marie for days at the Omniplex and evenings of hot chocolate; and for inspiring me to pursue my doctorate.

I would like to thank my grandparents Jesse and Eva for unconditional love and support. I would especially like to thank my grandma for letting my grandpa take me to Slippery Falls all those years.

I would like to thank my brothers, sisters and friends for sharing their lives with me.

Finally, I would like to thank Sunny, my constant companion. I can't imagine my life without him.

Chapter 1

Equilibrium with finitely additive mixed strategies

1.1 Introduction

We consider two equilibrium concepts in normal form games in which we allow players access to strategies intended to model vague beliefs. Examples of the beliefs that we are trying to model include “I will play my first action some day” and “He will charge a price slightly higher than \$1.” On a technical level, this is achieved by requiring that players’ mixed strategies be only finitely additive instead of countably additive. We find that payoffs are ambiguously defined when more than one player chooses a vague strategy. This ambiguity presents difficulty when a player is considering a potential deviation from her current strategy. We identify two interesting ways of resolving this ambiguity, leading us to two equilibrium concepts.

The weaker equilibrium notion, which we call an *optimistic equilibrium*, may be interpreted as an equilibrium in which players are able to rationalize the candidate equilibrium strategy profile by resolving ambiguities in the most optimistic way possible when considering the candidate profile and resolving them in the most pessimistic way possible when considering deviations. We are able to show that every game admits at least one optimistic equilibrium. This is a weak equilibrium concept, but the Bertrand-Edgeworth duopoly pricing example in section 1.2 indicates that it can provide interesting economic insight. In section 1.3, we analyze a well-known game formulated by Sion and Wolfe [SW57] with compact strategy spaces and not-particularly-pathological utility functions that has no Nash equilibrium, but that does have an optimistic equilibrium.

In the stronger equilibrium notion, which we call a *vague equilibrium*, an player is unwilling to deviate from a candidate equilibrium profile even if he believes that ambiguities in his payoff will resolve badly for him at the candidate profile, but will resolve in his favor when he deviates. This notion is more closely related to Nash equilibrium. In fact, every vague equilibrium is the limit of a sequence of ϵ -Nash equilibria. Consequently, we might think of these as idealized equilibria that stand in for sequences of ϵ -equilibria that are strategically

stable as ϵ tends to 0. Stated another way, vague equilibria classify sequences of ϵ -equilibria (with ϵ tending to 0) whose “limits” are not equilibria simply because the limiting strategies do not exist rather than because of some interesting strategic discontinuity in the limit.

The work most closely related to ours is [Mar97], in which players are given access to finitely additive mixed strategies. However, [Mar97] considers only payoff functions satisfying a restriction that excludes a large number of economically interesting games.¹ When payoff functions satisfy the restriction, all of our equilibrium concepts coincide and our results are equivalent.

It is well known that the standard method for proving the existence of a mixed strategy equilibrium in a normal form game using the Glicksberg-Kakutani-Fan fixed point theorem does not work when the players’ pure strategy sets are not compact or the payoff functions are discontinuous. There has been substantial work in finding conditions that ensure the existence of equilibria in discontinuous games which is formally independent of our work, but from which we have benefitted greatly [BM13, Car09, DM86a, MMT11, Pro11, Ren99]. However, our goal is different. Rather than looking for restrictions on games that ensure the existence of an equilibrium, we seek equilibrium concepts that allow us to understand games which may not have equilibria in the usual sense.

The remainder of the chapter is structured as follows. We begin with a few informal examples of our equilibrium concepts in sections 1.2, 1.3, and 1.4. We summarize the mathematics needed for our theory in section 1.5 and present the game theoretic setup in section 1.6. We consider first the weaker of our two equilibrium concepts in section 1.7 and the stronger in section 1.8. In section 1.9, we consider various relationships between our equilibrium concepts and Nash equilibrium. In section 1.10, we consider a technical restriction on utility functions that make our extensions of the agents’ utility functions easier to work with. We conclude in section 1.11.

1.2 The Bertrand-Edgeworth model of duopoly pricing with limited supply

We consider a model of duopoly pricing for a homogenous good in which each seller has a restricted quantity of the good to sell. This is the model considered in [DM86b]. It is proved that there exists a mixed equilibrium in this game, but, to our knowledge, no explicit equilibrium is known.

There are two sellers and a single good. Seller i has a stock S_i of the good to sell. Each buyer is represented by a point on the unit interval and total demand for the good at price p is given by $D(p)$. Each seller chooses a price $p_i \geq 0$. The buyers will choose to purchase from the seller with the lowest price. Any unserved buyers may then purchase from the second seller if they wish. If the sellers charge the same price, they capture the share of the market in proportion to their stock until their stock is exceeded.

¹see section 1.6.

In summary, the payoff function for the first seller (and analogously for the second seller) is

$$u_1(p_1, p_2) = \begin{cases} \min\{p_1 S_1, p_1 D(p_1)\} & \text{if } p_1 < p_2 \\ \min\left\{p_1 S_1, p_1 D(p_1) \frac{S_1}{S_1 + S_2}\right\} & \text{if } p_1 = p_2 \text{ and } S_2 \geq \frac{D(p_1)S_2}{S_1 + S_2} \\ p_1(D(p_1) - S_2) & \text{if } p_1 = p_2 \text{ and } S_2 < \frac{D(p_2)S_2}{S_1 + S_2} \\ \max\left\{0, p_1 D(p_1) \frac{D(p_2) - S_2}{D(p_2)}\right\} & \text{if } p_1 > p_2. \end{cases}$$

We will make a few assumptions that are not made in [DM86b] to simplify our analysis. We will assume that D is continuous, strictly decreasing and that there exist \underline{p} and \bar{p} for which $D(\underline{p}) = S_1 + S_2$ and $D(\bar{p}) = 0$. We will also assume that the unrestricted monopoly profit function $\Pi(p) = pD(p)$ is strictly concave on $[\underline{p}, \bar{p}]$.

We will write p^- for the vague strategy in which the seller plans to charge a price infinitesimally smaller than p .² If both sellers choose to play p^- , then there is some ambiguity about which is first-to-market. If each seller resolves the ambiguity in the most optimistic way possible, then each will believe that she will be first-to-market and that her price will be indistinguishable from p . That is, seller i believes that she will recognize the payoff $\min\{pS_i, \Pi(p)\}$ from the profile (p^-, p^-) .

Theorem 1. *There are numbers $a, b \in \mathbb{R}$ defined in the proof below such that when $p \in [a, b]$, neither seller will have incentive to deviate from the profile (p^-, p^-) when they resolve their payoff ambiguity optimistically.³*

Proof. Let $\pi_1(p) = \min\{pS_1, \Pi(p)\}$. The function π_1 is seller 1's first-to-market payoff function. That is, $\pi_1(p)$ is the profit that seller 1 would achieve given that the other seller chooses a price higher than p . Let p_1^* be the unique point in $[\underline{p}, \bar{p}]$ at which π_1 achieves its maximum. The existence of this point is guaranteed by our assumptions on Π . Notice that seller 1's second-to-market profit is her first to market profit scaled down by the factor $(D(p_2) - S_2)/D(p_2)$. Notice also that the sellers will never choose to match prices exactly. Consequently, if seller 2 chooses a price $p_2 < p_1^*$, the price p_1^* will dominate all other prices that do not make seller 1 first-to-market. That is, seller 1 will respond either with a price that makes her first-to-market or with the price p_1^* .

Define $\hat{p}_1 \in (\underline{p}, p_1^*)$ by the equation

$$\pi_1(\hat{p}_1) = \pi_1(p_1^*)(D(\hat{p}_1) - S_2)/D(\hat{p}_1).$$

Our assumptions on D and Π guarantee that \hat{p}_1 is well defined. This definition of \hat{p}_1 ensures that

$$\pi_1(p) \geq p_1^* D(p_1^*)(D(p) - S_2)/D(p)$$

²In the formalism of the sequel, p^- will stand for a finitely additive 0-1 measure that assigns probability 1 to every interval of the form $(p - \epsilon, p)$ with $\epsilon > 0$. See our remarks at the end of section 1.6.

³In general, it is not necessarily true that $a \leq b$, in which case the theorem is vacuously true. However, we are guaranteed that $a \leq b$ when the game is symmetric, so the theorem is not vacuous in general.

as long as $\hat{p}_1 \leq p \leq p_1^*$. That is, seller 1 has no incentive to deviate to a higher price from any price $p \in (\hat{p}_1, p_1^*)$ if she believes that she is first-to-market at the price p and that her deviation will leave her second-to-market facing the price p on the part of seller 2.

We see that if seller 1 believes that she is first to the market with a price $p \in [\underline{p}, p_1^*]$, she will have no incentive to lower her price since π_1 is increasing on $[\underline{p}, p_1^*]$ and will have no incentive to increase her price if it means that she will be second to the market.

We may repeat the same analysis for seller 2. Define $a = \max\{\hat{p}_1, \hat{p}_2\}$ and $b = \min\{p_1^*, p_2^*\}$. (In the symmetric case, $\hat{p}_1 = \hat{p}_2 \leq p_2^* = p_1^*$, so $a \leq b$.) Then, neither seller will have any incentive to deviate from the profile (p^-, p^-) when $p \in [a, b]$.

□

This furnishes our first example of an *optimistic equilibrium*. Informally, this means that the players are optimistic about their payoffs at the equilibrium profile, but are pessimistic about their payoffs when they consider deviations. We will make this concept precise in section 1.7. Notice that the strategy profile (p, p) , which closely resembles (p^-, p^-) , is not an ϵ -Nash equilibrium for ϵ sufficiently close to 0.

We propose the following interpretation of this equilibrium. As long as prices fall within a certain range, sellers will tend to set similar prices, but each will try to gain an infinitesimal advantage so as to be first-to-market. This result is highly suggestive of a phenomenon known as Edgeworth price cycling [KRRS94]. An Edgeworth price cycle is a dynamic description of prices in which, as long as the prices fall in some range, we observe a price war and prices decrease together. As soon as prices are low enough that one seller is better off being second-to-market, that seller increases her price. This is quickly followed by a price increase by the other seller and the cycle starts anew.

Our optimistic equilibrium concept provides some justification for each seller's behavior during an Edgeworth price cycle. The sellers' pricing decisions at any point in time are quite reasonable provided each seller believes that she will get the advantage in the price war.

1.3 A game with no value

Consider the following two-player, zero-sum game introduced by Sion and Wolfe [SW57]. Each player's action set is $[0, 1]$ and the payoff to player 1 is

$$u_1(x, y) = \begin{cases} 1 & \text{if } x > y \text{ or } x + \frac{1}{2} < y \\ 0 & \text{if } x = y \text{ or } x + \frac{1}{2} = y \\ -1 & \text{if } x < y < x + \frac{1}{2}. \end{cases}$$

It is shown in [SW57] that this game has no equilibrium in countably additive mixed strategies.

We allow each player access to strategies of the form x^+ and x^- where x^+ means some infinitesimal amount larger than x and x^- means some infinitesimal amount smaller than x . We may encounter payoff ambiguities when both players consider a strategy of this type.




Figure 1.1: A game with no value

For example, $1, 0$ and -1 are all plausible payoffs at the strategy profile $(1/2^-, 1/2^-)$. We will write $\bar{u}_1(x, y)$ for the highest plausible payoff at the strategy profile (x, y) and $\underline{u}_1(x, y)$ for the lowest. So, $\bar{u}_1(1/2^-, 1/2^-) = 1$ and $\underline{u}_1(1/2^-, 1/2^-) = -1$.

We will say that the strategy x_1 dominates x_2 if $\underline{u}_1(x_1, y) \geq \underline{u}_1(x_2, y)$ for all strategies y . That is, we compare the worst-case outcomes when deciding that one strategy dominates another.

We will find an equilibrium by first eliminating dominated strategies. Notice that the strategy $x = 1$ on the part of player 1 dominates w for each $1/2 \leq w < 1$ (including strategies like $1/2^+$). We see that $\underline{u}_1(1, y) = 1$ if $y \neq 1$ and $\underline{u}_1(1, 1) = 0$. Moreover, $\underline{u}_1(w, 1) \leq 0$ for all $w \in [1/2, 1)$. Player 1's remaining strategies are then $[0, 1/2) \cup \{1\}$.

Next, the strategy $y = 1/2^-$ of player 2 dominates each strategy z with $0 \leq z < 1/2^-$. Player 2's relevant payoff at $y = 1/2^-$ is

$$\underline{u}_2(w, 1/2^-) = \begin{cases} 1 & \text{if } 0 \leq w < 1/2^- \\ -1 & \text{if } w = 1/2, 1/2^-, 1. \end{cases}$$

Each strategy $y \in [0, 1/2^-)$ does as poorly as $1/2^-$ in the worst case against $w = 1/2, 1/2^-, 1$. Player 2's remaining strategies are then $[1/2^-, 1]$.

Now, the strategy $x = 0$ on the part of player 1 dominates each strategy w with $0 < w < 1/2$. Player 1's payoff from $x = 0$ is

$$\underline{u}_1(0, z) = \begin{cases} -1 & \text{if } z = 1/2^- \\ 0 & \text{if } z = 1/2 \\ 1 & \text{if } z > 1/2. \end{cases}$$

However, every strategy $w \in (0, 1/2)$ has a worst-case payoff of -1 against $z = 1/2^-$ or $z = 1/2$. Notice that using the worst-case evaluation is necessary here to eliminate $w = 1/2^-$. Player 1 is then left with the strategy set $\{0, 1\}$.

Then, each strategy $y \in (1/2, 1)$ returns a payoff of -1 against each remaining strategy of player 1, so we may eliminate these. Player 2's strategy set is then $\{1/2^-, 1/2, 1\}$.

Then, player 2's strategy $y = 1/2^-$ does better than $z = 1/2$ in the worst case against $w = 0$. Both $y = 1/2^-$ and $z = 1/2$ perform equally poorly against $w = 1$ in the worst case, so we may eliminate $z = 1/2$. Player 2 is then left with the strategies $\{1/2^-, 1\}$.

We are left with the game with payoff matrix:

$$\begin{array}{c|cc} & \frac{1}{2}^- & 1 \\ \hline 0 & -1 & 1 \\ \hline 1 & 1 & 0 \end{array}$$

It is straightforward to show that player 1 playing $x = 0$ with probability $1/3$ and $x = 1$ with probability $2/3$ and player 2 playing $y = 1/2^-$ with probability $1/3$ and $y = 1$ with probability $2/3$ is an equilibrium with payoff $1/3$. This furnishes our second example of an optimistic equilibrium.

1.4 The Bertrand competition model of duopoly pricing with anti-predatory pricing regulation

In this model, two sellers are able to supply an unlimited quantity of a good at a fixed per unit cost of c . The total market demand at a price p is $D(p)$, which we will assume bounded and decreasing. The payoff to seller 1 (and analogously for seller 2) is

$$u_1(p_1, p_2) = \begin{cases} (p_1 - c)D(p_1) & \text{if } p_1 < p_2 \\ \frac{1}{2}(p_1 - c)D(p_1) & \text{if } p_1 = p_2 \\ 0 & \text{if } p_1 > p_2. \end{cases}$$

As stated so far, most of our readers will recognize that the price profile (c, c) is the unique Nash equilibrium in this game. However, suppose that there is also an anti-predatory pricing law that prevents the duopolists from selling at or below cost.⁴

Let us allow the duopolists access to strategies of the form p^+ and p^- as before. Suppose that the duopolists are playing the profile (c^+, c^+) . If each duopolist believes that they will receive the smallest possible payment at this profile, then each believes that she will make a profit of 0. However, every other price will result in a profit of 0, even if any ambiguities are resolved as optimistically as possible. The profile (c^+, c^+) furnishes our first example of a vague equilibrium. That is, this equilibrium suggests that players will charge prices very close to cost, but will not charge exactly at cost to avoid prosecution.

⁴This is reasonably close to the interpretation of the Sherman Antitrust Act as argued in [AT75], although setting a price *exactly* equal to the cost would not be prohibited.

Notice that although there is a unique vague equilibrium, we still have optimistic equilibria of the form (p^-, p^-) for every $p > c$ for which $p' \mapsto p'D(p')$ is non-decreasing on $[c, p]$. That is, in this model we also recover seemingly collusive strategy profiles as optimistic equilibria. This could provide a non-collusive justification for instances in which prices fail to fall instantaneously to the competitive level after a competitor enters a monopolist's market.

1.5 Mathematical preliminaries

Let (X, Σ) be a measure space with X a separable metric space and Σ the Borel σ -algebra.

Definition 1. We will say that $\mu : \Sigma \rightarrow \mathbb{R}$ is a bounded, finitely additive measure if

(i) there is some $M > 0$ such that $|\mu(E)| < M$ for all $E \in \Sigma$;

(ii) $\mu\left(\bigcup_{j=1}^K E_j\right) = \sum_{j=1}^K \mu(E_j)$ for every finite, pairwise disjoint collection E_1, \dots, E_K from Σ .

Furthermore, we will say that μ is countably additive if the last condition holds for every countable, pairwise disjoint collection $\{E_j\}_{j=1}^\infty$.

The word measure without qualification will mean a bounded, finitely additive measure. The set of bounded, measurable functions on X will be denoted $\mathcal{F}(X)$. The set of bounded, finitely additive measures on X will be denoted $\text{ba}(X)$. The weak* topology on $\text{ba}(X)$ is the smallest (coarsest) topology that makes the map

$$\sigma \mapsto \int_X f d\sigma$$

continuous for each $f \in \mathcal{F}(X)$.

Definition 2. We will call a measure $\mu \in \text{ba}(X)$ a probability measure if

(i) $\mu(X) = 1$

(ii) $\mu(A) \geq 0$ for all $A \in \Sigma$.

Theorem 2 (Banach-Alaoglu). *The set of probability measures on X is weak* compact.*

Proof. See [AB06], theorem 6.25, for example. □

It is the Banach-Alaoglu theorem that makes the space of finitely additive probability measures useful as an extension of the space of countably additive probability measures. Instead of thinking about extending the space of countably additive probability measures on X to the space of finitely additive probability measures, we could instead think about expanding X to a larger space and then considering countably additive measures on that

space. A remarkable theorem of Yosida and Hewitt shows that these perspectives are equivalent [YH52], as we will now see.

Let $\Omega(X)$ be the set of finitely additive measures ω on X that take only the values 0 and 1⁵ (i.e., $\omega(A) = 0$ or $\omega(A) = 1$ for all $A \in \Sigma$). For any $x \in X$, the point mass δ_x is in $\Omega(X)$, so we may regard X as a subset of $\Omega(X)$. We have assumed that X is a separable metric space precisely so that every countably additive 0-1 measure is a point mass.

We will regard $\Omega(X)$ as a topological space with the weak* topology that it inherits as a subset of $\text{ba}(X)$. We have a description of this topology that is perhaps more intuitive. Given $A \in \Sigma$, define

$$\bar{A} = \{\omega \in \Omega(X) : \omega(A) = 1\}.$$

Each set \bar{A} is open and sets of the form \bar{A} generate the weak* topology on $\Omega(X)$.

Theorem 3. $\Omega(X)$ is compact.

Proof. This follows immediately from the Banach-Alaoglu theorem and from the fact that $\Omega(X)$ is closed in the unit ball in $\text{ba}(X)$. \square

Theorem 4. X is a dense open subset of $\Omega(X)$.

Proof. The weak* topology is generated by sets of the form

$$\begin{aligned} W(\tau, \{B_j\}_{j=1}^k, \{\epsilon_j\}_{j=1}^k) &= \{\sigma \in \Omega(X) : |1_{B_j}(\sigma) - 1_{B_j}(\tau)| < \epsilon_j, j = 1, \dots, k\} \\ &= \{\sigma \in \Omega(X) : \sigma(B_j) = \tau(B_j), j = 1, \dots, k\} \end{aligned}$$

with $0 < \epsilon_j < 1$ and B_j measurable for all j . Replacing B_j by B_j^c as necessary, we may assume that $\tau(B_j) = 1$ for all j . From the formula

$$\tau(C \cap D) = \tau(C) + \tau(D) - \tau(C \cup D),$$

we see that the intersection of any two sets of τ -measure 1 must have measure 1. It follows that $\bigcap_{j=1}^k B_j$ is non-empty. Then,

$$\delta_x \in W(\tau, \{B_j\}_{j=1}^k, \{\epsilon_j\}_{j=1}^k)$$

for any $x \in \bigcap_{j=1}^k B_j$.

To see that X is open, note that $\overline{\{x\}} = \{x\}$ and $X = \bigcup_{x \in X} \{x\}$. \square

Given a bounded, measurable function $f : X \rightarrow \mathbb{R}$, we may extend it to a function on $\Omega(X)$ by defining $f : \Omega(X) \rightarrow \mathbb{R}$ by

$$f(\omega) = \int_X f d\omega.$$

⁵Some readers may prefer to think of these as ultrafilters on Σ .

Notice that $f(\delta_x) = f(x)$ for all $x \in X$. By definition of the weak* topology, the map $\omega \mapsto f(\omega)$ is continuous.

It turns out that every probability measure⁶ on X may be represented as a *countably additive* measure on $\Omega(X)$.

Theorem 5 (Yosida-Hewitt [YH52], comment 4.5). *Let σ be a probability measure on X . There is a regular, countably additive Borel probability measure $\bar{\sigma}$ on $\Omega(X)$ such that for every bounded, measurable function f on X ,*

$$\int_X f d\sigma = \int_{\Omega(X)} f(\omega) d\bar{\sigma}(\omega) = \int_{\Omega(X)} \int_X f(x) d\omega(x) d\bar{\sigma}(\omega).$$

We also know that every probability measure may be approximated by a linear combination of 0-1 measures.

Theorem 6 (Yosida-Hewitt [YH52], theorem 4.6). *Every probability measure γ is the weak* limit of linear combinations of 0-1 measures, in the sense that for every $\epsilon > 0$ and every collection $\{f_j\}_{j=1}^\ell$ of bounded, measurable functions, there is a measure $\delta = \sum_{i=1}^k \alpha_i \omega_i$ with each $\omega_i \in \Omega(X)$ and each α_i a scalar such that*

$$\left| \int_X f_j d\gamma - \int_X f_j d\delta \right| < \epsilon$$

for all $j = 1, \dots, \ell$.

We will state a few results that are not necessary for the sequel, but that are interesting in their own right and may be helpful for understanding the space of finitely additive probability measures. Let $\bar{\Sigma}$ be the σ -algebra generated by sets of the form \bar{A} for $A \in \Sigma$. In general $\bar{\Sigma}$ is a sub- σ -algebra of the set of Borel measurable subsets of $\Omega(X)$.

Lemma 1. *If $A \in \bar{\Sigma}$, then $A \cap X \in \Sigma$.*

Proof. Let Λ be the collection of sets A in $\bar{\Sigma}$ for which $A \cap X \in \Sigma$. Λ contains each set of the form \bar{B} with $B \in \Sigma$ since $\bar{B} \cap X = B$. Moreover, it is a σ -algebra, so $\Lambda = \bar{\Sigma}$. \square

We have a characterization of countably additive probability measures, which we believe we are the first to observe.

Theorem 7. *Let σ be a probability measure on X . Then, σ is countably additive if and only if $\bar{\sigma}(A) = \sigma(A \cap X)$ for all $A \in \bar{\Sigma}$.*

⁶We will state our results for probability measures since they are all that we need in the sequel. However, Yosida and Hewitt show in theorem 1.12 [YH52] that every measure may be written as a difference of nonnegative measures, which may be used to generalize the following results to all of $\text{ba}(X)$.

Proof. Suppose first that σ is countably additive. Let D be the collection of sets A in $\bar{\Sigma}$ for which $\bar{\sigma}(A) = \sigma(A \cap X)$. We know that $\bar{A} \in D$ for each $A \in \bar{\Sigma}$ regardless of the countable additivity of σ since $\bar{A} \cap X = A$. Our goal then is to show that D is a Dynkin system. We know already that $\Omega(X) \in D$. Suppose that $A, B \in D$ with $A \subseteq B$. Then,

$$\begin{aligned}\bar{\sigma}(B \setminus A) &= \bar{\sigma}(B) - \bar{\sigma}(A) \\ &= \sigma(B \cap X) - \sigma(A \cap X) \\ &= \sigma((B \setminus A) \cap X),\end{aligned}$$

so $B \setminus A \in D$. Suppose now that $A_1, A_2, \dots \in D$ with $A_n \subseteq A_{n+1}$. Then,

$$\begin{aligned}\bar{\sigma}\left(\bigcup_{n=1}^{\infty} A_n\right) &= \lim_{n \rightarrow \infty} \bar{\sigma}(A_n) \\ &= \lim_{n \rightarrow \infty} \sigma(A_n \cap X) \\ &= \sigma\left(\bigcup_{n=1}^{\infty} (A_n \cap X)\right) \\ &= \sigma\left(\left(\bigcup_{n=1}^{\infty} A_n\right) \cap X\right),\end{aligned}$$

where we used the countable additivity of σ in the second to last equality. It follows that $\bigcup_{n=1}^{\infty} A_n \in D$, so D is a Dynkin system and the result follows.

Suppose now that $\bar{\sigma}(A) = \sigma(A \cap X)$ for all $A \in \bar{\sigma}$. Suppose that $A_n \downarrow \emptyset$ with $A_n \in \Sigma$. Then,

$$\begin{aligned}\lim_{n \rightarrow \infty} \sigma(A_n) &= \lim_{n \rightarrow \infty} \bar{\sigma}(\bar{A}_n) \\ &= \bar{\sigma}\left(\bigcap_{n=1}^{\infty} \bar{A}_n\right) \\ &= \sigma\left(\bigcap_{n=1}^{\infty} \bar{A}_n \cap X\right) \\ &= \sigma\left(\bigcap_{n=1}^{\infty} A_n\right) \\ &= \sigma(\emptyset) \\ &= 0.\end{aligned}$$

□

Yosida and Hewitt show that every probability measure may be written as a sum of a countably additive measure and a measure which is purely finitely additive in the following sense.

Definition 3. We will say that a probability measure σ is purely finitely additive if the only countably additive measure τ for which $0 \leq \tau \leq \sigma$ is $\tau = 0$.

Theorem 8 (Yosida-Hewitt [YH52], theorem 1.23). *Let σ be any probability measure. Then, σ may be written uniquely as the sum of a countably additive measure σ_c and a purely finitely additive measure σ_p .*

We are able to show that purely finitely additive probability measures are confined to a closed, nowhere dense $G_{\delta\sigma}$.

Theorem 9. *Let $\sigma \geq 0$ be a purely finitely additive probability measure on X . There is a set $\mathcal{B} \in \overline{\Sigma}$ such that $\mathcal{B} \cap X = \emptyset$ and $\overline{\sigma}(\mathcal{B}) = \overline{\sigma}(\Omega(X))$. Moreover, we may choose \mathcal{B} to be a $G_{\delta\sigma}$.⁷*

This is a weak version of the converse to the following theorem of Yosida and Hewitt.

Theorem 10 (Yosida-Hewitt [YH52], theorem 4.16). *Let σ be a probability measure. If $\overline{\sigma}$ is confined to a closed, nowhere dense G_{δ} ,⁸ then σ is purely finitely additive.*

Under the hypotheses of Yosida and Hewitt's theorem, we may find a sequence $A_1 \supset A_2 \supset \dots$ of sets from Σ such that $\bigcap_{n=1}^{\infty} A_n = \emptyset$, but $\bigcap_{n=1}^{\infty} \overline{A_n}$ is the closed G_{δ} on which σ is concentrated. That is, $\lim_{n \rightarrow \infty} \sigma(A_n) = 1$.

It is reasonable to conjecture that the converse is true. That is, if σ is purely finitely additive, then there is a sequence $A_1 \supset A_2 \supset \dots$ of sets from σ such that $\bigcap_{n=1}^{\infty} A_n = \emptyset$, but $\lim_{n \rightarrow \infty} \sigma(A_n) = 1$. That is, if σ is purely finitely additive, there is a single decreasing sequence of sets that demonstrates that fact. Neither proof nor counterexample for this conjecture has been forthcoming. Our theorem tells us that while there may not be a single decreasing sequence of sets, we can find a countable number of sequences decreasing to nothing that carry all of the mass of σ .

Proof of theorem 9. We will call $\{A_n\}_{n=1}^{\infty}$ from Σ a null sequence if $A_n \downarrow \emptyset$. Let

$$\alpha = \sup \left\{ \lim_{n \rightarrow \infty} \sigma(A_n) : \{A_n\}_{n=1}^{\infty} \text{ is a null sequence} \right\}.$$

For each $k \in \mathbb{N}$, choose a null sequence $\{A_n^k\}_{n=1}^{\infty}$ for which $\lim_{n \rightarrow \infty} \sigma(A_n^k) \geq \alpha - 1/k$.

Define $B_n^k = \bigcup_{j=1}^k A_n^j$. We see that for each k , $\{B_n^k\}_{n=1}^{\infty}$ is a null sequence for which $\lim_{n \rightarrow \infty} \sigma(B_n^k) \geq \alpha - 1/k$. Moreover, $B_n^k \subseteq B_n^{k+1}$ for all n, k .

Define $\mathcal{B}^k = \bigcap_{n=1}^{\infty} \overline{B_n^k}$. We see that $\mathcal{B}^k \subseteq \mathcal{B}^{k+1}$. Moreover, $\overline{\sigma}(\mathcal{B}^k) \geq \alpha - 1/k$. Let $\mathcal{B} = \bigcup_{k=1}^{\infty} \mathcal{B}^k$. We see that $\overline{\sigma}(\mathcal{B}) = \alpha$.

Define the measure μ on X by $\mu(A) = \overline{\sigma}(A \cap \mathcal{B}^c)$. Since $0 \leq \mu \leq \sigma$, if μ is countably additive then we are done. Suppose then that $\mu \neq 0$. Then, there is some $\epsilon > 0$ and a null sequence $\{C_n\}_{n=1}^{\infty}$ for which $\lim_{n \rightarrow \infty} \mu(C_n) = \epsilon$. Letting $\mathcal{C} = \bigcap_{n=1}^{\infty} \overline{C_n}$, we see that $\overline{\sigma}(\mathcal{C} \cap \mathcal{B}^c) = \epsilon$.

⁷A $G_{\delta\sigma}$ is a countable union of countable intersections of open sets.

⁸A G_{δ} is a countable intersection of open sets.

Now, for each k , $\{B_n^k \cup C_n\}_{n=1}^\infty$ is a null sequence. Since $\{B_n^k\}_{n=1}^\infty$ and $\{C_n\}_{n=1}^\infty$ are decreasing, we have

$$\bigcap_{n=1}^\infty (B_n^k \cup C_n) = \left(\bigcap_{n=1}^\infty B_n^k \right) \cup \left(\bigcap_{n=1}^\infty C_n \right) = \mathcal{B}^k \cup \mathcal{C}.$$

So,

$$\begin{aligned} \lim_{n \rightarrow \infty} \sigma(B_n^k \cup C_n) &= \bar{\sigma}(\mathcal{B}^k \cup \mathcal{C}) \\ &\geq \bar{\sigma}(\mathcal{B}^k) + \bar{\sigma}(\mathcal{C} \cap \mathcal{B}^c) \\ &\geq \alpha - 1/k + \epsilon. \end{aligned}$$

As soon as $1/k < \epsilon$, we find a null sequence that contradicts our choice of α . It follows that $\mu = 0$ and the result follows. \square

Corollary 1. *Let σ be a purely finitely additive probability measure. Then, $\bar{\sigma}(X) = 0$.*

Proof. $\bar{\sigma}(\mathcal{B}^c) = 0$ and $X \subseteq \mathcal{B}^c$. \square

It would be very tidy if it were true that $\bar{\sigma}(X^c) = 0$ for every countably additive σ , but alas, that is not the case as our next result shows.

Theorem 11. *Suppose that σ is a finitely additive probability measure on X and that $\bar{\sigma}(X) = 1$. Then, σ is countably additive and atomic.*

Proof. It follows from corollary 1 that σ is countably additive. Since $\bar{\sigma}$ is regular,

$$\bar{\sigma}(X) = \sup\{\bar{\sigma}(K) : K \text{ compact, } K \subseteq X\}.$$

Suppose that K is an infinite compact subset of X . Then, we may choose a sequence $\{x_n\}_{n=1}^\infty$ of distinct elements of K . This sequence has a convergent subnet, which cannot be an element of X since it assigns probability zero to each point in the sequence and to each singleton in X not in the sequence. It follows that no such K may exist. So, for each n , we may choose a finite set K_n for which $\bar{\sigma}(K_n) \geq 1 - 1/n$. Letting $A = \bigcup_{n=1}^\infty K_n$, we see that A is countable and $\bar{\sigma}(A) = 1$. \square

Definition 4. *We will say that a probability measure σ is atomic on an algebra \mathcal{A} containing the singletons if there is a sequence of points $\{a_k\}_{k=1}^\infty$ in A such that $\lim_{k \rightarrow \infty} \sigma\{a_1, \dots, a_k\} = 1$. We will say that a point a_k of the sequence is an atom if $\sigma\{a_k\} > 0$.*

The only subtlety is that $\{a_1, a_2, \dots\}$ need not be \mathcal{A} -measurable. In what follows, our next theorem will guarantee that profiles of atomic probability measures will give us a unique product probability measure.

Theorem 12. *Let \mathcal{A} be an algebra containing the singletons that generates Σ as a σ -algebra. Suppose that σ is a probability measure on Σ such that $\sigma|_{\mathcal{A}}$ is atomic. Then, σ is countably additive and atomic.*

Proof. Let the sequence $\{a_k\}_{k=1}^\infty$ be as in the definition of an atomic probability measure. Let $A_1 \supset A_2 \supset A_3 \cdots$ be any sequence of sets in Σ decreasing to the empty set and let $\epsilon > 0$. Pick K large enough that $\sigma\{a_1, \dots, a_K\}^c < \epsilon$. For large enough N , we must have $A_N \subset \{p_1, \dots, p_K\}^c$, which implies that $\sigma(A_N) < \epsilon$. It follows that σ is countably additive. \square

1.6 Game theoretic setup

We consider a normal form game with N players. Each player has a strategy space A_i , which we assume to be a separable metric space. Let \mathcal{A}_i be the Borel σ -algebra on A_i , let $A = \prod_{i=1}^N A_i$, and let \mathcal{A} be the algebra on A generated by the measurable boxes. We will assume that each payoff function $u_i : A \rightarrow \mathbb{R}$ is measurable with respect to $\sigma(\mathcal{A})$, the smallest σ -algebra containing \mathcal{A} .

Let $P_i = \Omega(A_i)$, the space of 0-1 measures on A_i , and let $P = \prod_{i=1}^N P_i$. Finally, let F_i be the space of countably additive probability measures on P_i and let $F = \prod_{i=1}^N F_i$. We will call strategies in F_i *vague strategies*. Theorem 5 tells us that F_i is the same as the space of finitely additive probability measures on A_i . We would like to extend each u_i to a weak*-continuous function on F , but this is not generally possible.

Consider for example the zero-sum game in which each player chooses a natural number and the player who selects the larger number receives payment of 1. Let ν be any weak* limit of the sequence $\{\delta_n\}_{n \in \mathbb{N}}$. Let us try to figure out what $u_1(\nu, \nu)$ would be if u_1 extended continuously to $\Omega(\mathbb{N}) \times \Omega(\mathbb{N})$. On the one hand, we would have

$$u_1(\nu, \nu) = \lim_{n \rightarrow \infty} \lim_{m \rightarrow \infty} u_1(\delta_n, \delta_m) = -1.$$

On the other hand, we would have

$$u_1(\nu, \nu) = \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} u_1(\delta_n, \delta_m) = 1.$$

We see that, in general, our utility functions need not extend continuously to our expanded strategy spaces.

The only general condition on u_i that we know of that ensures that $u_i : A \rightarrow \mathbb{R}$ extends to a continuous function $u_i : F \rightarrow \mathbb{R}$ is that u_i be \mathcal{A} -measurable [Mar97]. Since this restriction excludes a large number of interesting games, we will focus instead on upper- and lower-semicontinuous extensions of the players' payoff functions.

We will do this in two ways. First, we will define a dense subset S of F on which we know that our utility functions are well defined. We will then take the least upper-semicontinuous extension (denoted \bar{U}_i) and largest lower-semicontinuous extension (denoted \underline{U}_i) to all of F . Second, since A is a dense subset of P , we may extend each utility function to the least upper-semicontinuous extension (denoted \bar{u}_i) and the largest lower-semicontinuous extension (denoted \underline{u}_i) on P . We can then extend \bar{u}_i and \underline{u}_i to functions on F by integration.

Let S be the space of profiles of linear combinations of point masses. That is, $\sigma \in S$ if $\sigma_i = \sum_{k=1}^{\ell_i} \lambda_i^k \delta_{a_i^k}$ with $a_i^k \in A_i$ and $\sum_{k=1}^{\ell_i} \lambda_i^k = 1$ for some $\ell_i < \infty$ for each i . Since

finite sets of points in A are \mathcal{A} -measurable, we see that each profile in S specifies a unique probability measure on A . Moreover, this probability measure is countably additive. These “well-behaved” profiles are dense in the space of all profiles.

Lemma 2. S is dense in F .

Proof. This follows from theorems 4 and 6. □

Let Γ_i be the graph of $U_i : S \rightarrow \mathbb{R}$ regarded as a subset of $F \times \mathbb{R}$ and let $\bar{\Gamma}_i$ be the closure of Γ_i . Define $\bar{U}_i : F \rightarrow \mathbb{R}$ by

$$\bar{U}_i(\sigma) = \inf\{u : (\sigma, u) \in \bar{\Gamma}_i\}.$$

Similarly, $\underline{U}_i : F \rightarrow \mathbb{R}$ is defined by

$$\underline{U}_i(\sigma) = \inf\{u : (\sigma, u) \in \bar{\Gamma}_i\}.$$

This gives us our first set of extensions of u_i .

For the second set of extensions, let Δ_i be the graph of $u_i : A \rightarrow \mathbb{R}$ regarded as a subset of $P \times \mathbb{R}$ and let $\bar{\Delta}_i$ be the closure of Δ_i . Define

$$\bar{u}_i(p) = \sup\{u : (p, u) \in \bar{\Delta}_i\}$$

and similarly

$$\underline{u}_i(p) = \inf\{u : (p, u) \in \bar{\Delta}_i\}.$$

Then, define $\bar{V}_i : F \rightarrow \mathbb{R}$ and $\underline{V}_i : F \rightarrow \mathbb{R}$ by

$$\bar{V}_i(\sigma_1, \dots, \sigma_N) = \int_P \bar{u}_i d(\bar{\sigma}_1 \otimes \dots \otimes \bar{\sigma}_N)$$

and

$$\underline{V}_i(\sigma_1, \dots, \sigma_N) = \int_P \underline{u}_i d(\bar{\sigma}_1 \otimes \dots \otimes \bar{\sigma}_N).$$

Since $\bar{\sigma}_i$ is countably additive on P_i for each i , there is no issue with defining the product measure. As a result, we will write $\bar{\sigma}$ for $\bar{\sigma}_1 \otimes \dots \otimes \bar{\sigma}_N$ and similarly for $\bar{\sigma}_{-i}$.

Notice that the map \bar{V}_i is upper-semicontinuous and \underline{V}_i is lower-semicontinuous since the weak* topology on F_i corresponds to the topology of weak convergence on the space of countably additive measures on P_i . The functions \bar{V}_i and \underline{V}_i are the second set of extensions of u_i to F .

We can say a few things about the relationship between our various ways of extending our utility functions.

Theorem 13. For all $\sigma \in F$ and all i , $\underline{V}_i(\sigma) \leq \underline{U}_i(\sigma) \leq \bar{U}_i(\sigma) \leq \bar{V}_i(\sigma)$.

Proof. We know that \bar{V}_i and \bar{U}_i are upper-semicontinuous and \underline{V}_i and \underline{U}_i are lower-semicontinuous. Moreover, \bar{U}_i is the smallest upper-semicontinuous extension of U_i to F and \underline{U}_i is the largest lower-semicontinuous extension of U_i to F . □

All of our extensions agree when all of the players choose pure strategies and at most one player takes advantage of a non-countably additive strategy.

Theorem 14. *Let $p \in P$ such that at most one p_i is not countably additive. That is, there is some j such that for each $i \neq j$, there exists $a_i \in A_i$ such that $p_i = \delta_{a_i}$. Then,*

$$\bar{V}_i(p) = \underline{V}_i(p) = \int_{A_j} U_i(a_1, \dots, a_N) dp_j(a_j).$$

Proof. Without loss of generality, assume $j = 1$. Suppose that we have a net

$$(\delta_{a_1^\alpha}, \dots, \delta_{a_N^\alpha}, U_i(a_1^\alpha, \dots, a_N^\alpha)) \rightarrow (p_1, \delta_{a_2}, \dots, \delta_{a_N}, u)$$

for some $u \in \mathbb{R}$.

We may choose $\beta \in A$ so that $\delta_{\alpha_i^N}(1_{a_i}) = 1$ for all $i \geq 2$ for all $\alpha \geq \beta$. This implies that

$$\begin{aligned} (\delta_{a_1^\alpha}, \dots, \delta_{a_N^\alpha}, U_i(a_1^\alpha, \dots, a_N^\alpha)) &= (\delta_{a_1^\alpha}, \delta_{a_2}, \dots, \delta_{a_N}, U_i(a_1^\alpha, \dots, a_N)) \\ &= \left(\delta_{a_1^\alpha}, \delta_{a_2}, \dots, \delta_{a_N}, \int_{A_1} U_i(a_1, \dots, a_N) d\delta_{a_1^\alpha}(a_1) \right) \\ &\rightarrow \left(p_1, \delta_{a_2}, \dots, \delta_{a_N}, \int_{A_1} U_i(a_1, \dots, a_N) dp_1(a_1) \right) \end{aligned}$$

for all $\alpha \geq \beta$. It follows that

$$u = \int_{A_1} U_i(a_1, \dots, a_N) dp_1(a_1). \quad \square$$

This extends to countably additive atomic strategies.

Theorem 15. *Suppose that $\sigma \in F$ such that σ_i is countably additive and atomic for all but at most one i . Then, $\underline{V}_i(\sigma) = \bar{V}_i(\sigma)$. This implies that $\underline{U}_i(\sigma) = \bar{U}_i(\sigma)$.*

Proof. Let j be the player for which σ_j is not necessarily countably additive and atomic. It follows from theorem 12 that σ_{-j} specifies a unique probability measure on A_{-j} and that this probability measure is countably additive. Let $T_{-j} \subseteq A_{-j}$ be the set of atoms of the measure σ_{-j} . Then,

$$\begin{aligned} \underline{V}_i(\sigma) &= \int_P \underline{u}_i(p) d(\bar{\sigma}_1 \otimes \dots \otimes \bar{\sigma}_N)(p) \\ &= \sum_{a_{-j} \in T_{-j}} \int_{P_j} \underline{u}_i(p_j, a_{-j}) d\bar{\sigma}_j(p_j) \\ &= \sum_{a_{-j} \in T_{-j}} \int_{P_j} \bar{u}_i(p_j, a_{-j}) d\bar{\sigma}_j(p_j) \\ &= \int_P \bar{u}_i(p) d\bar{\sigma}(p) \\ &= \bar{V}_i(\sigma) \end{aligned} \quad \square$$

The atomic assumption may appear strange to some readers. We conjecture that it may be safely removed, but have been unable to prove that this is the case. In section 1.10, we prove that there is a large class of utility functions, which we call *agreeable functions*, for which we may remove the atomic assumption. All of the utility functions in examples that we consider in this paper fall into this class.

It turns out that the upper-semicontinuous extensions agree and the lower-semicontinuous extensions agree when restricted to P , although the upper-semicontinuous extensions need not agree with the lower-semicontinuous extensions.

Theorem 16. $\bar{U}_i(p) = \bar{u}_i(p)$ and $\underline{U}_i(p) = \underline{u}_i(p)$ for all $p \in P$.

Proof. In view of theorem 13, it suffices to show that $\bar{u}_i(p) \leq \bar{U}_i(p)$ and similarly for the lower-semicontinuous extensions. This follows immediately from the following facts. First, theorem 14 implies that all of the extensions agree on A . In particular, \bar{U}_i extends $u_i : A \rightarrow \mathbb{R}$. That is, the restriction of \bar{U}_i to P is an upper-semicontinuous extension of u_i . Since \bar{u}_i is the *smallest* upper-semicontinuous extension of $u_i : A \rightarrow \mathbb{R}$ to P , the result follows. \square

In the event that the strategy spaces are one-dimensional it is often convenient to work with certain equivalence classes of elements of P . Suppose that each A_i is a Borel-measurable subset of \mathbb{R} . For any $a_i \in A_i$ we will write a_i^- for the set of measures in P_i that assign probability 1 to each of the sets $(a - \epsilon, a) \cap A_i$ with $\epsilon > 0$. Similarly, a_i^+ will denote the set of measures in P_i that assign probability 1 to each of the sets $(a, a + \epsilon) \cap A_i$ with $\epsilon > 0$. We will also abuse notation to write a_i for the set containing the probability measure δ_{a_i} . The sets $-\infty^+$ and ∞^- are defined analogously.

Our next theorem tells us that these sets partition A_i , so we may regard them as specifying an equivalence relation on A_i .

Lemma 3. *Suppose that A_i is a Borel-measurable subset of \mathbb{R} . Then, each element of p_i is a member of a_i^+, a_i^- , or a_i for some $a_i \in A_i \cup \{-\infty, \infty\}$.*

Proof. Let F be the cumulative distribution function of the probability measure $p_i \in P_i$. F is non-decreasing and takes (at most) the values 0 and 1. Let

$$a = \inf\{x : F(x) = 1\}.$$

If $\{x : F(x) = 1\}$ is empty, we set $a = \infty$ and see that p_i must assign probability 1 to each set of the form (M, ∞) with $M \in \mathbb{R}$, so $p_i \in \infty^-$. Similarly, if

$$a = -\infty,$$

we see that $p_i \in -\infty^+$. Suppose then that $a \in \mathbb{R}$. Let $\epsilon > 0$. Then,

$$p_i(a - \epsilon, a + \epsilon) = 1.$$

Written another way we see that

$$p_i(a - \epsilon, a) + p_i(\{a\}) + p_i(a, a + \epsilon) = 1.$$

Since p_i is a 0-1 measure, exactly one of these terms is 1 and the others are 0. Since the nonzero term must be the same for all $\epsilon > 0$, the result follows. \square

In our introductory examples, we were evaluating our utility functions at equivalence classes by saying, for example,

$$\bar{u}_i(a_1^+, a_2^-) = \sup_{p_1 \in a_1^+, p_2 \in a_2^-} \bar{u}_i(p_1, p_2).$$

Using these equivalence classes greatly simplifies working with finitely additive measures in one dimensional games, but caution is necessary as there is no guarantee that the supremum will be achieved simultaneously by the *same* pair (p_1, p_2) for both \bar{u}_1 and \bar{u}_2 . In each of our examples this problem may be solved by picking a single finitely additive probability measure 0^+ . We then interpret a^+ as the translation of this probability measure by a (i.e. $a^+(B) = 0^+(B - a)$), 0^- as the reflection of this measure across the origin, and a^- as the translation of 0^- by a . However, this cannot justify use of these equivalence classes in every situation.

1.7 Optimistic equilibria

We turn now to formalizing and investigating an equilibrium in which players evaluate the equilibrium profile favorably and deviations unfavorably.

Definition 5. *We will say that $\sigma \in F$ is an optimistic equilibrium if for each player i and each potential deviation $\tau_i \in F_i$,*

$$\bar{V}_i(\sigma) \geq \underline{V}_i(\tau_i, \sigma_{-i}).$$

Our most important result is that optimistic equilibria always exist.

Theorem 17. *Every game has an optimistic additive equilibrium.*

The proof follows the standard script. We will define a best response correspondence and observe that a fixed point corresponds to an optimistic equilibrium. We will then show that the best response correspondence satisfies the hypotheses of the Kakutani-Glicksberg-Fan fixed point theorem [Fan52, Gli52, K⁺41].

Definition 6. *We will say that a correspondence $\Gamma : X \rightarrow 2^X$ is Kakutani if each of the following hold*

- (i) X is a subset of a topological vector space;
- (ii) Γ is upper-hemicontinuous;
- (iii) $\Gamma(x)$ is non-empty, compact and convex for all $x \in X$.

Notice that if X is compact, to show that Γ is Kakutani, it suffices to show that Γ is upper-hemicontinuous and that $\Gamma(x)$ is non-empty and convex for all $x \in X$. Notice also that if $\{\Gamma^\alpha\}_{\alpha \in A}$ is any collection of Kakutani correspondences on a compact space X , then to show that $\bigcap_{\alpha \in A} \Gamma^\alpha$ is Kakutani, it suffices to show that $\bigcap_{\alpha \in A} \Gamma^\alpha(x)$ is non-empty for all $x \in X$.

Theorem 18 (Kakutani-Glicksberg-Fan). *Let X be a non-empty, convex, compact subset of a topological vector space. If $\Gamma : X \rightarrow 2^X$ is Kakutani, then Γ has a fixed point.*

We are now prepared to give the equilibrium existence proof.

Proof of theorem 17. For any $\gamma_i \in F_i$, define the correspondence

$$\mu \mapsto \text{BR}_i(\mu, \gamma_i) = \{\tau \in F : \bar{V}_i(\tau_i, \mu_{-i}) \geq \underline{V}_i(\gamma_i, \mu_{-i})\}.$$

We claim that the correspondence $\mu \mapsto \text{BR}_i(\mu, \gamma_i)$ is Kakutani. We will first show that this correspondence has a closed graph. Consider nets $\mu^\alpha \rightarrow \mu$ and $\tau^\alpha \rightarrow \tau$ with $\tau^\alpha \in \text{BR}_i(\mu^\alpha, \gamma_i)$. Then, since \bar{V}_i is upper-semicontinuous and \underline{V}_i is lower-semicontinuous,

$$\bar{V}_i(\tau_i, \mu_{-i}) \geq \limsup_{\alpha} \bar{V}_i(\tau_i^\alpha, \mu_{-i}^\alpha) \geq \liminf_{\alpha} \underline{V}_i(\gamma_i, \mu_{-i}^\alpha) \geq \underline{V}_i(\gamma_i, \mu_{-i}).$$

That is, $\tau \in \text{BR}_i(\mu, \gamma_i)$. Since F is compact, this implies that $\mu \mapsto \text{BR}_i(\mu, \gamma_i)$ is upper-hemicontinuous.

To see that $\text{BR}_i(\mu, \gamma_i)$ is non-empty, note that $(\gamma_i, \mu_{-i}) \in \text{BR}_i(\mu, \gamma_i)$.

To show that $\text{BR}_i(\mu, \gamma_i)$ is convex, suppose that $\sigma, \tau \in \text{BR}_i(\mu, \gamma_i)$ and $\lambda \in (0, 1)$. We have

$$\bar{V}_i(\lambda\sigma_i + (1 - \lambda)\tau_i, \mu_{-i}) = \lambda\bar{V}_i(\sigma_i, \mu_{-i}) + (1 - \lambda)\bar{V}_i(\tau_i, \mu_{-i}) \geq \underline{V}_i(\gamma_i, \mu_{-i}).$$

It follows that $\text{BR}_i(\mu, \gamma_i)$ is convex. We have shown that $\mu \mapsto \text{BR}_i(\mu, \gamma_i)$ is Kakutani.

Define a new correspondence by

$$\mu \mapsto \text{BR}_i(\mu) = \bigcap_{\gamma_i \in F_i} \text{BR}_i(\mu, \gamma_i).$$

To show that BR_i is Kakutani, we need only show that $\text{BR}_i(\mu)$ is non-empty. Since F is compact and $\text{BR}_i(\mu, \gamma_i)$ is compact for each γ_i , it suffices to show that $\bigcap_{\ell=1}^k \text{BR}_i(\mu, \gamma_i^\ell)$ is non-empty for every finite collection $\{\gamma_i^1, \dots, \gamma_i^k\} \subset F_i$. Choose γ_i^j such that $\bar{V}_i(\gamma_i^j) \geq \bar{V}_i(\gamma_i^\ell)$ for all $\ell = 1, \dots, k$. Then, $(\gamma_i^j, \mu_{-i}) \in \bigcap_{\ell=1}^k \text{BR}_i(\mu, \gamma_i^\ell)$. It follows that BR_i is Kakutani.

As before, to show that $\text{BR} = \bigcap_{i=1}^N \text{BR}_i$ is Kakutani, it suffices to show that $\text{BR}(\mu)$ is non-empty for all μ . This is clear from the definition of BR_i .

The result now follows from the Kakutani-Glicksberg-Fan theorem. \square

Having proved that optimistic equilibria always exist, we will prove a few results to help us identify them. First, optimistic equilibria have the single deviation property.

Theorem 19. *A profile $\sigma \in F$ is an optimistic equilibrium if and only if for every i and every $p_i \in P_i$, $\bar{V}_i(\sigma) \geq \underline{V}_i(p_i, \sigma_{-i})$.*

Proof. Suppose that τ_i is a favorable deviation for player i : $\bar{V}_i(\sigma) < \underline{V}_i(\tau_i, \sigma_{-i})$. Then, since

$$\underline{V}_i(\tau_i, \sigma_{-i}) = \int_{P_i} \int_{P_{-i}} \underline{u}_i(p_i, p_{-i}) d\bar{\sigma}_{-i}(p_{-i}) d\bar{\tau}_i(p_i).$$

there must be at least one p_i for which

$$\underline{V}_i(p_i, \sigma_{-i}) = \int_{P_{-i}} \underline{u}_i(p_i, p_{-i}) d\bar{\sigma}_{-i} > \bar{V}_i(\sigma).$$

The other implication is part of the definition of an optimistic equilibrium. \square

One of the most useful tools at our disposal for finding optimistic equilibria is iterated elimination of dominated strategies.

Definition 7. We will say that $p_i \in P_i$ worst-case dominates $q_i \in P_i$ if $\underline{u}_i(p_i, \gamma_{-i}) \geq \underline{u}_i(q_i, \gamma_{-i})$ for all $\gamma_{-i} \in P_{-i}$.

Of course, we need to show that iterated elimination of worst-case dominated strategies is legitimate in the sense that an optimistic equilibrium obtained after eliminating strategies is an optimistic equilibrium in the original game.

Theorem 20. Suppose that we remove some or all of the strategies worst-case dominated by some strategy $p_i \in P_i$ (other than p_i itself). If σ is an optimistic equilibrium in the resulting game then it is an optimistic equilibrium in the original game.

Proof. Let σ be an optimistic equilibrium after some collection of strategies worst-case dominated by p_i have been removed. Suppose that q_i is a potential deviation for player i that was eliminated. Then,

$$\begin{aligned} \underline{V}_i(q_i, \sigma_{-i}) &= \int_{P_{-i}} \underline{u}_i(q_i, p_{-i}) d\bar{\sigma}_{-i}(p_{-i}) \\ &\leq \int_{P_{-i}} \underline{u}_i(p_i, p_{-i}) d\bar{\sigma}_{-i}(p_{-i}) \\ &= \underline{V}_i(p_i, \sigma_{-i}) \\ &\leq \bar{V}_i(\sigma). \end{aligned}$$

The result then follows from theorem 19. \square

1.8 Vague equilibria

We now turn to an equilibrium notion which is very strong in the sense that the equilibrium profile is evaluated unfavorably and potential deviations are evaluated favorably.

Definition 8. A strategy profile $\tau \in F$ is a vague equilibrium if, for each player i and all $\gamma_i \in F_i \setminus \{\tau_i\}$,

$$\underline{U}_i(\tau) \geq \bar{U}_i(\gamma_i, \tau_{-i}).$$

We call this vague equilibrium because it is, in our view, the most reasonable equilibrium notion when players have access to vague strategies. Analogously with optimistic equilibria, to show that a profile is a vague equilibrium it suffices to show that it is immune to pure deviations.

Lemma 4. A strategy profile $\sigma \in F$ is a vague equilibrium if and only if for every i and every $p_i \in P_i$, $\bar{U}_i(p_i, \sigma_{-i}) \leq \underline{U}_i(\sigma)$.

Proof. The forward implication is part of the definition of a vague equilibrium.

Suppose then that σ is not a vague equilibrium. Then, there is some $\tau_i \in F_i$ such that $\underline{U}_i(\sigma) < \bar{U}_i(\tau_i, \sigma_{-i})$. Pick $\{(\tau_i^\alpha, \sigma_{-i}^\alpha)\}_{\alpha \in A}$ from S such that $(\tau_i^\alpha, \sigma_{-i}^\alpha) \rightarrow (\tau_i, \sigma_{-i})$ and $U_i(\tau_i^\alpha, \sigma_{-i}^\alpha) \rightarrow \bar{U}_i(\tau_i, \sigma_{-i})$. For each α , there is some $a_i^\alpha \in A_i$ for which $U_i(a_i^\alpha, \sigma_{-i}^\alpha) \geq U_i(\tau_i^\alpha, \sigma_{-i}^\alpha)$. Pass to a subnet on which $a_i^\alpha \rightarrow p_i$ for some $p_i \in P_i$. Then,

$$\bar{U}_i(p_i, \sigma_{-i}) \geq \limsup_{\alpha} U_i(a_i^\alpha, \sigma_{-i}^\alpha) \geq \bar{U}_i(\tau_i, \sigma_{-i}) > \underline{U}_i(\sigma). \quad \square$$

Our only real tool for finding vague equilibria is iterated elimination of dominated strategies.

Definition 9. We will say that a strategy $p_i \in P_i$ decisively dominates $q_i \in P_i$ if $\underline{u}_i(p_i, \gamma_{-i}) \geq \bar{u}_i(q_i, \gamma_{-i})$ for all $\gamma_{-i} \in P_{-i}$.

Vague equilibria behave as we might expect with respect to iterated elimination of decisively dominated strategies. That is, if we are able to find a vague equilibrium in the game that results from iterated elimination of decisively dominated strategies, it will be a vague equilibrium in the original game.

Theorem 21. Suppose that we remove some or all of the strategies decisively dominated by some strategy p_i (other than p_i itself). If σ is a vague equilibrium after the removal of these strategies then it is a vague equilibrium with the strategies still in place.

Proof. Let σ be a vague equilibrium after the weakly dominated strategies have been eliminated. Suppose that player i has some potential deviation $q_i \in P_i$ that was eliminated.

Then,

$$\begin{aligned}
 \bar{U}_i(q_i, \sigma_{-i}) &\leq \bar{V}_i(q_i, \sigma_{-i}) \\
 &= \int_{P_{-i}} \bar{u}_i(q_i, p_{-i}) d\bar{\sigma}_{-i} \\
 &\leq \int_{P_{-i}} \underline{u}_i(p_i, p_{-i}) d\bar{\sigma}_{-i} \\
 &= \underline{V}_i(p_i, \sigma_{-i}) \\
 &\leq \bar{U}_i(p_i, \sigma_{-i}) \\
 &\leq \underline{U}_i(\sigma).
 \end{aligned}$$

The result then follows from lemma 4. □

1.9 Relations between equilibrium concepts

We see that vague equilibrium is a stronger notion than optimistic equilibrium.

Theorem 22. *Suppose that σ is a vague equilibrium. Then, σ is an optimistic equilibrium.*

Proof. This is immediate from the definitions. □

We will use Nash equilibrium to mean exactly what it usually means. It is a fixed point of the best response correspondence when players have access to *countably additive* mixed strategies. It turns out that vague equilibrium is stronger than Nash equilibrium when it makes sense to compare them.

Theorem 23. *Suppose that σ is a vague equilibrium and that σ is countably additive. Then, σ is a Nash equilibrium.*

Proof. Any profitable deviation in the countably additive extension would give us a deviation in the finitely additive extension. □

While we expect for the analogous result to be true for optimistic equilibria, we have been unable to prove it. The issue is that we are unable to show that all of our utility function extensions agree when evaluated at a not-necessarily-atomic, countably additive mixed strategy. Instead, we have the following weaker result that uses the theory of agreeable functions developed in section 1.10.

Theorem 24. *Suppose that σ is an optimistic equilibrium and that σ is countably additive. Suppose also that each player's utility function is agreeable. Then, σ is a Nash equilibrium.*

Proof. The proof is the same as that of the previous theorem in view of theorem 27 below. □

Our next result suggests another way of finding optimistic equilibria. They may appear as limits of ϵ -Nash equilibria as ϵ tends to 0.

Theorem 25. *Suppose that $\{\sigma^n\}_{n=1}^\infty$ is a sequence of countably additive strategy profiles such that σ^n is an ϵ^n -Nash equilibrium. Then, any limit point of $\{\sigma^n\}_{n=1}^\infty$ is an optimistic equilibrium.*

Proof. Let $\{\sigma^\alpha\}_{\alpha \in A}$ be a subnet of $\{\sigma^n\}_{n=1}^\infty$ that converges to σ . Let $p_i \in P_i$ be some potential deviation for player i . Pick a net $\{a_i^\alpha\}_{\alpha \in A}$ from A_i for which $a_i^\alpha \rightarrow p_i$.⁹ Then,

$$\begin{aligned} \underline{U}_i(p_i, \sigma_{-i}) &\leq \liminf_{\alpha \in A} U_i(a_i^\alpha, \sigma_{-i}^\alpha) \\ &\leq \liminf_{\alpha \in A} (U_i(\sigma^\alpha) + \epsilon^\alpha) \\ &\leq \bar{U}_i(\sigma^\alpha). \end{aligned} \quad \square$$

We might hope that a sequence of ϵ -Nash equilibria would converge to the stronger vague equilibrium, but, as the following example shows, that is not the case.

Example 1. Consider the two-player, symmetric game in which player 1 chooses some $x_1 \in (0, 1)$ and then receives payoff x_1 if $x_1 \neq x_2$ and 0 if $x_1 = x_2$. For any $\epsilon > 0$, the strategy profile $(1 - \epsilon/2, 1 - \epsilon/4)$ is an ϵ -Nash equilibrium. Now, we have

$$\lim_{\epsilon \rightarrow 0} U_1(1 - \epsilon/2, 1 - \epsilon/4) = 1$$

and similarly for player 2. However, the profile $(1^-, 1^-)$ is not a vague equilibrium. The worst case payoff is obtained via the sequence $(1 - \epsilon, 1 - \epsilon)$. This sequence implies that

$$\underline{U}_1(1^-, 1^-) \leq \lim_{\epsilon \rightarrow 0} U_1(1 - \epsilon, 1 - \epsilon) = 0.$$

Since $\bar{U}_1(1/2, 1^-) = 1/2 > 0$, the strategy $1/2$ is a profitable deviation for player 1.

While it is not true that every limit of ϵ -equilibria is a vague equilibrium, every vague equilibrium is the limit of ϵ -equilibria. Moreover, these ϵ -equilibria are very simple in the sense that each player plays a linear combination of point masses.

Theorem 26. *Suppose that σ is a vague equilibrium. Then, there is a net $\{(\sigma^\alpha, \epsilon^\alpha)\}_{\alpha \in A}$ such that $\sigma^\alpha \rightarrow \sigma$ and $\epsilon^\alpha \rightarrow 0$ and σ^α is an ϵ^α -Nash equilibrium.*

Proof. Let $\{\sigma^\alpha\}_{\alpha \in A}$ be a net of profiles in S that converges to σ . Let ϵ^α be the smallest ϵ for which σ^α is an ϵ -Nash equilibrium. Let $\epsilon = \liminf_{\alpha \in A} \epsilon^\alpha$ and pass to a subnet for which $\epsilon^\alpha \rightarrow \epsilon$. If $\epsilon = 0$, then we are done. Otherwise, assume $\epsilon > 0$. For each α , there is some

⁹A priori, this net may have a different index set than $\{\sigma^\alpha\}_{\alpha \in A}$, but the product construction for nets that allows us to find a single net $\{(\sigma^\beta, p_i^\beta)\}_{\beta \in B}$ for which $(\sigma^\beta, p_i^\beta) \rightarrow (\sigma, p_i)$.

player i^α who has an $\epsilon^\alpha/2$ -favorable deviation. Pass to a subnet on which $i^\alpha = i$ for some i . Then,

$$\begin{aligned} \underline{U}_i(\sigma) &\leq \liminf_{\alpha \in A} U_i(\sigma^\alpha) \\ &< \liminf_{\alpha \in A} \left(U_i(a_i^\alpha, \sigma_{-i}^\alpha) - \frac{\epsilon^\alpha}{2} \right) \\ &\leq \bar{U}_i(p_i, \sigma_{-i}) - \frac{\epsilon}{2}, \end{aligned}$$

but we know that $\underline{U}_i(\sigma) \geq \bar{U}_i(p_i, \sigma_{-i})$. This contradiction shows that $\epsilon = 0$. □

1.10 Agreeable functions

We now consider a condition which we will call agreeability on functions which ensures that, as far as the agreeable functions are concerned, every profile of countably additive strategies extends *uniquely* to a probability measure on A . We start by defining agreeability for measurable sets. Recall that \mathcal{A} is the *algebra* generated by the boxes in A and $\sigma(\mathcal{A})$ is the σ -algebra generated by \mathcal{A} .

Definition 10. *We will say that a set $B \in \sigma(\mathcal{A})$ is agreeable if for every profile $(\sigma_1, \dots, \sigma_N)$ of countably additive strategies and every finitely additive measure τ on $\sigma(\mathcal{A})$ that agrees with $\sigma_1 \otimes \dots \otimes \sigma_N$ on \mathcal{A} , we have $\tau(B) = (\sigma_1 \otimes \dots \otimes \sigma_N)(B)$.*

First, it is clear that the collection of agreeable sets is an algebra extending \mathcal{A} . We will call this algebra \mathcal{G} .

Definition 11. *We will say that a function $g : A \rightarrow \mathbb{R}$ is agreeable if it is \mathcal{G} -measurable.*

We care about agreeable functions because we know that if τ is any finitely additive probability measure that agrees with $\sigma_1 \otimes \dots \otimes \sigma_N$ with each σ_i countably additive and if f is any agreeable function, then

$$\int_A f d\tau = \int_A f d(\sigma_1 \otimes \dots \otimes \sigma_N).$$

In the context of extending our utility functions to the profiles of finitely additive measures, this implies the following.

Theorem 27. *If u_i is agreeable, then $\underline{U}_i(\sigma) = \bar{U}_i(\sigma)$ for all countably additive σ .*

We want to have quite a few agreeable sets. First, it is immediate that every \mathcal{A} -measurable function is agreeable. We have a few ways of finding \mathcal{A} -measurable functions.

Proposition (Marinacci [Mar97]). *Let $f_i : A_i \rightarrow \mathbb{R}$ be \mathcal{A}_i -measurable for each i . Then,*

(i) $F(x_1, \dots, x_N) = \prod_{i=1}^N f_i(x_i)$ is \mathcal{A} -measurable;

(ii) $F(x_1, \dots, x_N) = \sum_{i=1}^N f_i(x_i)$ is \mathcal{A} -measurable.

Proposition (Marinacci [Mar97]). *Let f be \mathcal{A} -measurable and suppose that g is a continuous function on \mathbb{R} . Then, $g \circ f$ is \mathcal{A} -measurable.*

We have the following proposition which (after proving theorems 28 and 29) will cover all of the utility functions that we are interested in.

Proposition. *Suppose that f^1, \dots, f^K are \mathcal{A} -measurable functions and G^1, \dots, G^K are \mathcal{G} -measurable sets. Then, $\sum_{k=1}^K f^k 1_{G^k}$ is agreeable.*

Proof. This is immediate from the definition of an agreeable function. \square

It may be the case that every $\sigma(\mathcal{A})$ -measurable function is agreeable. However, we are unable to show this.

Conjecture 1. *Every $\sigma(\mathcal{A})$ -measurable set is agreeable: $\sigma(\mathcal{A}) = \mathcal{G}$.*

Since we are unable to prove the conjecture, we will satisfy ourselves with a theorem that will ensure that the sets that we are interested in are agreeable. In particular, the graphs of measurable functions are agreeable as the corollary to the next theorem shows.

Theorem 28. *Let J be some proper subset of $\{1, \dots, N\}$ and let $f : A_{-J} \rightarrow A_J$ be any $(\mathcal{A}_{-J}, \mathcal{A}_J)$ -measurable function. Let D_f be the graph of f :*

$$D_f = \{(a_{-J}, f(a_{-J})) : a_{-J} \in A_{-J}\}.$$

Let σ be a profile of countably additive probability measures and let τ be any probability measure on A which agrees with σ on \mathcal{A} . Let Q_J be the atoms of σ_J . Then,

$$\tau(D_f) = \sum_{q \in Q_J} \sigma_{-J}(f^{-1}(\{q\})) \times \sigma_J(\{q\}).$$

Since the expression for $\tau(D_f)$ depends only on the value of τ on \mathcal{A} , we obtain the following corollary.

Corollary 2. *D_f is agreeable.*

Proof of theorem 28. Let $0 < \epsilon < 1$. For each $j \in J$, let Q_j^ϵ be a finite collection of atoms (possibly empty) such that σ_j assigns probability less than ϵ to the atoms not in Q_j^ϵ .

First, we see that $f^{-1}(\{q\}) \times \{q\} \subseteq D_f$ when $q \in Q_J^\epsilon = \prod_{j \in J} Q_j^\epsilon$. It follows that

$$\tau(D_f) \geq \sum_{q \in Q_J^\epsilon} \sigma_{-J}(f^{-1}(\{q\})) \times \sigma_J(\{q\}).$$

Letting $\epsilon \rightarrow 0$, we have $Q_J^\epsilon \uparrow Q_J$, from which we conclude that

$$\tau(D_f) \geq \sum_{q \in Q_J} \sigma_{-J}(f^{-1}(\{q\})) \times \sigma_J(\{q\}).$$

Next, let \mathcal{B}_j^ϵ be a finite collection of pairwise disjoint sets from \mathcal{A}_j such that

$$A_j = Q_j^\epsilon \cup \bigcup_{B \in \mathcal{B}_j^\epsilon} B$$

and for all $B \in \mathcal{B}_j^\epsilon$, B and Q_j^ϵ are disjoint and $\sigma_j(B) < \epsilon$. Let \mathcal{C}_J^ϵ be the collection of sets of the form $\prod_{j \in J} C_j$ with either $C_j \in \mathcal{B}_j^\epsilon$ or C_j an atom of σ_j . Notice that $\sigma_J(C_J) < \epsilon$ unless each C_j is an atom. We have

$$D_f \subseteq \bigcup_{C_J \in \mathcal{C}_J} f^{-1}(C_J) \times C_J.$$

By definition of C_J and f , we see that each of these sets is \mathcal{A} -measurable. It follows that

$$\begin{aligned} \tau(D_f) &\leq \tau\left(\bigcup_{C_J \in \mathcal{C}_J} f^{-1}(C_J) \times C_J\right) \\ &= \sum_{C_J \in \mathcal{C}_J} \sigma_{-J}(f^{-1}(C_J)) \cdot \sigma_J(C_J) \\ &\leq \sum_{q \in Q_j^\epsilon} \sigma_{-J}(f^{-1}(\{q\}) \times \sigma_J(\{q\}) + \sum_{C_J \in \mathcal{C}_J \setminus Q_j^\epsilon} \sigma_{-J}(f^{-1}(C_J)) \cdot \epsilon \\ &\leq \sum_{q \in Q_j^\epsilon} \sigma_{-J}(f^{-1}(\{q\})) \cdot \sigma_J(\{q\}) + \epsilon \\ &\leq \sum_{q \in Q_j} \sigma_{-J}(f^{-1}(\{q\})) \cdot \sigma_J(\{q\}) + \epsilon. \end{aligned}$$

Letting $\epsilon \rightarrow 0$, the result follows. \square

Theorem 29. *Suppose that B_1, \dots, B_k are open, pairwise disjoint subsets of A such that $\left(\bigcup_{j=1}^k B_j\right)^c$ is agreeable. Then, B_j is agreeable for each j .*

Proof. Let σ be a profile of countably additive probability measures and let τ be any probability measure on A which agrees with σ on \mathcal{A} . Let $C = \left(\bigcup_{j=1}^k B_j\right)^c$. We know that $\tau(C) = \sigma(C)$.

Since each B_j is open, we may write $B_j = \bigcup_{n=1}^\infty D_j^n$ with $D_j^n \in \mathcal{A}$. Moreover, we may choose the D_j^n such that $D_j^1 \subseteq D_j^2 \subseteq D_j^3 \subseteq \dots$.

Since $\tau(B_j) \geq \tau(D_j^n)$ for all n , we see that

$$\tau(B_j) \geq \lim_{n \rightarrow \infty} \tau(D_j^n) = \lim_{n \rightarrow \infty} \sigma(D_j^n) = \sigma(B_j).$$

We then have

$$\begin{aligned} 1 &= \tau(C) + \sum_{j=1}^k \tau(B_j) \\ &\geq \sigma(C) + \sum_{j=1}^k \sigma(B_j) \\ &= 1. \end{aligned}$$

It follows that $\tau(B_j) = \sigma(B_j)$ for each j . □

1.11 Conclusion

In this chapter, we have introduced two equilibrium concepts that are applicable when we allow players access to finitely additive strategies: optimistic equilibria and vague equilibria. Optimistic equilibria have the important property that they always exist. This provides, at the very least, a starting point in the analysis of any game. On a case-by-case basis, we may decide that an optimistic equilibrium are interesting economically, or, as may be the case even with Nash equilibria, we may decide that an optimistic equilibrium is simply an artifact of the process of formalizing an economic situation. In the Bertrand-Edgeworth duopoly example, we showed that optimistic equilibria can identify economically interesting strategy profiles even in games that are known to possess Nash equilibria. In our example, the optimistic equilibrium had the interpretation of justifying an instantaneous decision on the part of a duopolist in a dynamic setting.

Vague equilibria are more reasonable as a solution concept in the sense that they are closer to Nash equilibria. However, for this reason they tend not to identify new interesting strategy profiles. We regard their introduction as a contribution to the game theory vocabulary that allows us to make precise intuitive notions about the kinds of qualitative behavior that we might expect to see in normal form games that lack equilibria for compactness—rather than for strategic—reasons.

Chapter 2

Utility of wealth with many indivisibilities

2.1 Introduction

We investigate the properties of the utility of wealth function of an agent who chooses an optimal set of items from among a large number of indivisible items. Each item has a cost and provides some amount of utility to the agent. We are interested in the utility of wealth function obtained by solving this optimization problem for each wealth level. In the computer science literature this optimization problem is called the knapsack problem [KPP], so we will refer to the resulting utility function as a *knapsack utility function*.

It has been known since the seminal paper of Friedman and Savage [FS48] that an agent who seeks to maximize her expected utility may simultaneously gamble and purchase insurance when her utility function has a region of convexity sandwiched between regions of concavity. There has been a substantial amount of work in finding economic conditions that give rise to utility of wealth functions with convexities [AK81, Dob88, Hak70, HH13, Jon88, Kwa65, McC94]. Among these, Jones [Jon88], Kwang [Kwa65], and McCaffery [McC94] have suggested that an indivisibility in the consumption set may induce a region of convexity in the utility of wealth function and from this they recover Friedman and Savage's result that gambling may be part of an optimal utility maximization strategy.

We extend the results of Jones and Kwang to the situation in which all of the consumption goods are indivisible. By considering a model in which there are a large number of indivisibilities, we are able to consider the effect of an agent's wealth on the incentives to gamble that are caused by indivisibilities. In the single-indivisibility models presented by Jones and Kwang, if the agent is wealthy enough that she is past the region of convexity induced by the indivisibility she will prefer not to gamble. However, this conclusion appears to be an artifact of the assumption that there is a single indivisibility. In our model, we find that wealthy agents will tend to see relatively small (but sometimes positive) increases in expected utility from gambling as a result of large scale decreasing marginal utility.

By incorporating a large number of indivisibilities, we find that most wealth levels fall

in a region of convexity sandwiched between regions of concavity. As a result, we find that simultaneous gambling and insurance purchase is commonplace. This suggests that trying to predict the agent's behavior with respect to some gamble or contingent liability from a classification of the agent as "risk-loving" or "risk-averse" will likely be unsuccessful. Instead, an understanding of the main items relevant to the agent's consumption decision are necessary for good prediction. While the agent's attitude toward any particular risk depends a great deal on the particulars of the risk, we find that large gambling expenditures and a large monetary value placed on gambling will tend to result from the presence of high-cost, high-utility items that the agent is close to being able to afford.

We consider the applicability of knapsack utility functions to consumer behavior by considering the utility of wealth function in a continuous-time intertemporal model in which the agent's utility of consumption function at each point in time is a knapsack utility function. If the agent is able to borrow freely, the convexities which drive the interesting behavior of knapsack utility functions are absent from the utility of wealth function [Jon88]. However, after introducing a credit constraint we find that the agent's utility of wealth function converges to a knapsack utility function as the agent's saving rate becomes small. That is, the behavior of a credit-constrained agent with a low saving rate may be closer to that predicted by a knapsack utility function than by a concave utility function. To our knowledge, we are the first to use this model to justify transferring convexities in the utility of consumption function to the utility of wealth function.

In the intertemporal model, repeated negative-expected-return gambling may be rational in the sense that each gamble maximizes the agent's expected utility at the time that the agent undertakes it. From a societal perspective, this may be suboptimal because the law of large numbers implies that populations for whom gambling is rational will collectively become poorer. Moreover, members of these populations may find themselves in situations in which repeated participation in negative-expected-value gambles prevents their wealth from increasing with high probability. In short, our results suggest that indivisibility-induced gambling is a kind of poverty trap. In our conclusion, we point out the parameters of our model that could be targeted to reduce incentives for unfavorable gambling.

2.2 Knapsack utility functions

For us, an instantiation of the knapsack problem is specified by an infinite list of items, each with a cost and a utility. The lists $c = \{c_i\}_{i=1}^{\infty}$ and $u = \{u_i\}_{i=1}^{\infty}$ are the costs and utilities of the items. Given a wealth level w , a solution to the knapsack problem is a sequence $(a_i)_{i=1}^{\infty}$ of 0s and 1s that maximizes

$$\sum_{i=1}^{\infty} a_i u_i \quad \text{subject to} \quad \sum_{i=1}^{\infty} a_i c_i \leq w.$$

We may define the *utility of wealth function* by

$$U(w) = \sup_a \left\{ \sum_{i=1}^{\infty} a_i u_i : \sum_{i=1}^{\infty} a_i c_i \leq w \right\}$$

where the sup is taken over all lists $a = \{a_i\}_{i=1}^{\infty}$ with $a_i \in \{0, 1\}$.

We will make various assumptions about the collection of items. Define the *utility density* of item i to be $d_i = u_i/c_i$.

Assumptions. (i) For all i , $c_i > 0$ and $u_i \geq 0$.

(ii) $d_i \rightarrow 0$ and $d_1 > d_2 > \dots$.

(iii) There is some $C > 0$ for which $c_i \leq C$ for all i .

Assumption (i) states that the agent is a buyer rather than a seller and that all of the goods are positive goods.

Assumption (ii) states that the utility densities tend to 0. The assumption that $d_i \rightarrow 0$ implies that we may order our items so that $d_1 \geq d_2 \geq \dots$. This assumption is the analogue of the typical assumption of decreasing marginal utility of wealth. We insist that this sequence be strictly decreasing to make our results easier to state and prove.

Assumption (iii) asserts that the costs of items are bounded. If we regard indivisibilities as market imperfections that impede trade, we would expect for mechanisms to arise that divide these indivisible items, causing very large indivisibilities to be uncommon. We will not use this assumption most of the time. The primary importance of assumption iii is that it and assumption (ii) together imply that $u_i \rightarrow 0$.

Definition 12. We will call a function U a knapsack utility function if it is the utility of wealth function for a knapsack problem with item set satisfying assumptions (i)-(ii).

A related problem that is useful as a benchmark is the *linear relaxation* of the knapsack problem. In this problem, the constraint that we either buy or not buy a given item is relaxed so that we are allowed to buy any fraction of an item. The utility of wealth function in the linear relaxation of the knapsack problem is given by

$$\bar{U}(w) = \sup_a \left\{ \sum_{i=1}^{\infty} a_i u_i : \sum_{i=1}^{\infty} a_i c_i \leq w \right\},$$

where now the supremum is taken over all lists $a = \{a_i\}_{i=1}^{\infty}$ where $a_i \in [0, 1]$. The greedy algorithm is optimal for this problem. This is the algorithm that puts as much money as possible into the item with the highest utility density, then proceeds to the second highest utility density item and so on. We insist on assumption ii in the definition of a knapsack utility function precisely because it guarantees that the greedy algorithm is well-defined.

It follows that \bar{U} is concave. In fact, \bar{U} is the convex hull of the function U . That is, \bar{U} is the smallest concave function larger than U . Since $\bar{U}(w) \leq d_1 w$ for all w , we are able to conclude that $U(w) < \infty$ for all w . That is, our optimization problem never becomes infinite.

In figure 3.2, we show a knapsack utility function with its convex hull. This figure illustrates the features of knapsack utility functions that we find most interesting. First, the frequent oscillation between concavity and convexity can lead to concurrent gambling and

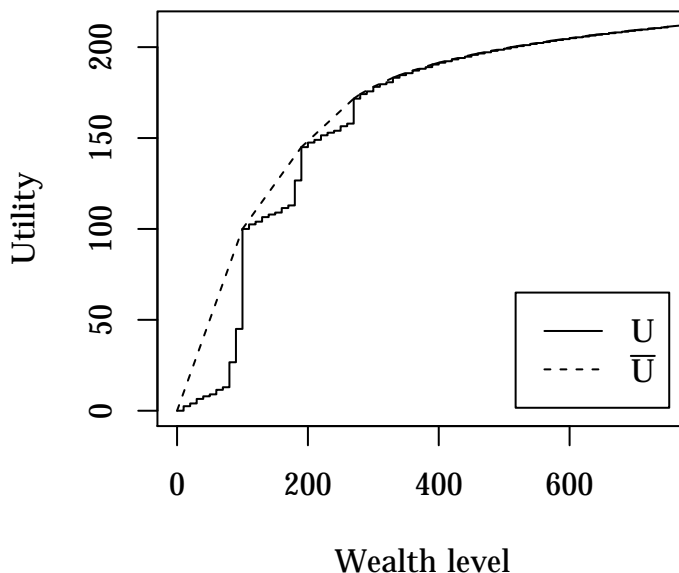


Figure 2.1: A knapsack utility function and its convex hull.

insurance purchase as in Friedman and Savage [FS48]. This is because the agent wants the next large purchase and is fearful of losing her last large purchase. Second, this oscillation is more pronounced at low wealth levels. This is caused by the fact that a wealthy agent will have already moved past the largest jumps in her utility function. That is, the agent will have already taken advantage of highest utility density opportunities.

2.3 Insurance and gambling

In this section, we will see that the concurrent purchase of insurance and gambles that motivated Friedman-Savage utility functions is not unusual when our agent uses a knapsack utility function. Our first result says that for most wealth levels, there is some contingent liability which the agent is willing to pay a premium to insure against. Our second result says that for most wealth levels, there is some gamble that the agent is willing to pay to accept. Putting the two results together, we find that the potential for concurrent purchase of gambling and insurance exists most of the time.

Definition 13. A contingent liability is a random variable L such that $L \leq 0$ and $E[L] < 0$. An insurance policy against the contingent liability L is any random variable of the form

$-L + E[L] - p$, where $p \geq 0$ is called the premium. A gamble is any random variable.¹

Let

$$W_i = c_1 + \cdots + c_i$$

for all $i = 1, 2, 3, \dots$ and let $W_0 = 0$. W_i is the wealth level at which the agent has purchased the i highest utility density items. Each W_i is important because $U(W_i) = \bar{U}(W_i)$, so gambling can only improve the agent's position if the gamble has positive expected value.

Theorem 30. *Suppose that the agent has wealth level $w \in (W_i, W_{i+1})$ for some $i = 1, 2, 3, \dots$ ² There is a contingent liability L that the agent is willing to pay a positive premium to insure against. Moreover, we may choose the contingent liability L and the premium p so that the agent's wealth after purchasing insurance lies in (W_i, W_{i+1}) .³*

The idea behind the proof of the theorem is that if the agent's wealth level falls below W_i , then the agent loses the i th highest utility density item.⁴ Since the items that the agent acquires between W_i and her current wealth w have relatively low utility density, she will be willing to sacrifice some of them to avoid losing the high utility density item.

Proof. Let $w' \in (W_{i-1}, W_i)$. Let $\alpha \in (0, 1)$ be such that

$$\alpha(w') + (1 - \alpha)(w) = W_i.$$

We know that

$$\alpha U(w') + (1 - \alpha)U(w) \leq \alpha \bar{U}(w') + (1 - \alpha)\bar{U}(w) < \bar{U}(W_i) = U(W_i).$$

By continuity of the map

$$\delta \mapsto \delta U(w') + (1 - \delta)U(w),$$

it follows that there is some $\epsilon > 0$ such that

$$\delta U(w') + (1 - \delta)U(w) < U(W_i)$$

for all $\delta \in (\alpha - \epsilon, \alpha)$.

Let $\delta \in (\alpha - \epsilon, \alpha)$ be fixed and let L be the contingent liability that costs $(w - w')$ with probability δ (and costs 0 otherwise). Our choice of δ allows us to find $p > 0$ such that $w + E[L] - p > W_i$.

¹We will use the convention that $U(w) = -\infty$ if $w < 0$ to avoid making any boundedness assumptions on the gambles.

²A similar result holds for almost all $w \in (W_0, W_1)$.

³All proofs may be found in the appendix.

⁴The agent's allocation at $w \in (W_i, W_{i+1})$ is difficult to compute exactly. It is likely that she will choose to purchase the i highest utility density items along with various lower utility density items, but this is not always the optimal allocation. When we say that she will "lose the i th highest utility density item" we mean that she will lose the jump in utility that she realized when her wealth reached W_i .

We see that the agent is willing to pay the premium p to insure against the contingent liability L :

$$\begin{aligned} U(w + E[L] - p) &> U(W_i) \\ &> \delta U(w') + (1 - \delta)U(w) \\ &= E[U(w + L)]. \end{aligned}$$

□

Theorem 31. *Suppose that the agent has wealth level $w \in (W_i, W_{i+1})$ for some $i = 0, 1, 2, \dots$. There is a gamble with negative expected value that the agent is willing to take.*

The idea behind the proof is that the agent is happy to accept the *mean-zero* gamble that leaves the agent with either wealth W_i or W_{i+1} . This gamble allows the agent to achieve expected utility equal to $\bar{U}(w) > U(w)$. Slightly increasing the probability with which the agent ends up with W_i leaves the agent's decision unchanged.

Proof of theorem 31. Let $\alpha \in (0, 1)$ such that

$$w = \alpha W_i + (1 - \alpha)W_{i+1}.$$

Let Z_δ be the gamble that costs $w - W_i$ with probability δ and pays W_{i+1} with probability $1 - \delta$. Since the sequence $(d_i)_{i=1}^\infty$ is strictly decreasing, the solution to the linear relaxation of the knapsack problem is unique. It follows that

$$E[U(w + Z_\alpha)] = \alpha U(W_i) + (1 - \alpha)U(W_{i+1}) = \bar{U}(w) > U(w).$$

By continuity of the map

$$\delta \mapsto E[U(w + Z_\delta)] = \delta U(W_i) + (1 - \delta)U(W_{i+1}),$$

there is some $\epsilon > 0$ such that

$$E[U(w + Z_\delta)] > U(w)$$

for all $\delta \in (\alpha, \alpha + \epsilon)$. The agent is willing to take any gamble Z_δ with $\delta \in (\alpha, \alpha + \epsilon)$ and these gambles all have negative expected value. □

These results are illustrated in figure 2.2.

2.4 Measures of risk seeking

The typical measures of risk-aversion and risk-seeking for differentiable utility functions are the Arrow-Pratt coefficients of absolute and relative risk-aversion.⁵ However, these measures are never well defined for knapsack utility functions as a result of the following theorem.

⁵Given a differentiable utility function u , the Arrow-Pratt coefficient of absolute risk-aversion is $-u''(w)/u'(w)$. The Arrow-Pratt coefficient of relative risk-aversion is $-wu''(w)/u'(w)$.




Figure 2.2: The agent begins with some non-random allocation, point A. Exposure to a contingent liability, moves her to point B. Purchasing an insurance policy with a positive premium moves her to point C. Finally, accepting a negative mean gamble moves her to point D. The agent would have obtained the same increase in expected utility by moving directly from B to D by simultaneously gambling and buying insurance.

Theorem 32. *The derivative of a knapsack utility function exists and is zero everywhere outside of a set of measure zero.*

Proof. This follows immediately from theorem 38, which we prove below. \square

We will investigate several alternative measures of risk-seeking behavior: the cost of the optimal gamble, the certainty equivalent of a gamble, and the change in expected utility from a gamble. For simplicity, we will consider only mean-zero gambles.

Definition 14. *The cost of a gamble Z which takes only finitely many values is the magnitude of the smallest value that Z takes with positive probability.*

If the cost of Z is c_Z , then we may write $Z = -c_Z + Z^+$ where $Z^+ \geq 0$ and $Z^+ = 0$ with positive probability. That is, the gamble Z provides the agent with the opportunity to pay c_Z for the non-negative random variable Z^+ .

Definition 15. *By an optimal gamble at wealth level w , we mean a mean-zero gamble Z^* for which*

$$E[U(w + Z^*)] \geq E[U(w + Z)]$$

for all mean-zero gambles Z .

We have an exact characterization of the optimal gambles.

Theorem 33. *Suppose that $w \in [W_i, W_{i+1})$ for some i . Then, the mean-zero gamble Z^* for which the agent's final wealth is either W_i or W_{i+1} is the unique optimal gamble.*

Proof. We compute

$$\begin{aligned} E[U(w + Z^*)] &= \frac{W_{i+1} - w}{W_{i+1} - W_i} \cdot U(w + W_i - w) + \frac{w - W_i}{W_{i+1} - W_i} \cdot U(w + W_{i+1} - w) \\ &= \frac{W_{i+1} - w}{W_{i+1} - W_i} \cdot U(W_i) + \frac{w - W_i}{W_{i+1} - W_i} \cdot U(W_{i+1}) \\ &= \bar{U}(w). \end{aligned}$$

The fact that Z^* is unique follow from the following two facts. If Z is mean-zero and takes values below the line ℓ connecting $(W_i, U(W_i))$ to $(W_{i+1}, U(W_{i+1}))$, then

$$E[U(w + Z)] < \bar{U}(w).$$

The line ℓ intersects the graph of U only at the points $(W_i, U(W_i))$ and $(W_{i+1}, U(W_{i+1}))$. \square

Heuristically, the agent's willingness to accept Z^* says that the agent is willing to gamble all of her items, except for the i highest utility density items, for the opportunity to have the $i + 1$ highest utility density items. Another way of thinking about Z^* is that it is the only gamble which allows the agent to achieve expected utility equal to $\bar{U}(w)$. Theorem 33 allows us to write down the cost of Z^* .

Corollary 3. *Let $w \in [W_i, W_{i+1})$ for some i . The cost of the optimal gamble at w is $w - W_i$, and is therefore bounded above by c_{i+1} .*

This characterizes the gambling behavior of an agent who may choose freely among mean-zero gambles. To the extent that our agent has free choice among gambles, large gambling expenditures are associated with large indivisibility costs.

We turn now to considering the certainty equivalent of a gamble.

Definition 16. *The certainty equivalent of a mean-zero gamble Z at the wealth level w is*

$$C(w, Z) = \sup\{b \in \mathbb{R} : U(w + b) \leq E[U(w + Z)]\}.$$

From the definitions of optimal gamble and certainty equivalent, we see that

$$C(w, Z^*) \geq C(w, Z)$$

for all mean-zero random variables Z . Since the certainty equivalent of the optimal gamble at $w \in (W_i, W_{i+1}]$ is bounded by $W_{i+1} - w$, we have the following corollary.

Corollary 4. *Let $w \in (W_i, W_{i+1}]$ for some i . For any mean-zero random variable Z , the certainty equivalent of Z at w is bounded above by $W_{i+1} - w$, and therefore by c_{i+1} .*

This continues the theme that large monetary values placed on gambling are associated with the presence of large, high utility density items. We will turn now to utility-based measures of risk-seeking behavior. We begin with a straightforward observation.

Theorem 34. *Let $w \in (W_i, W_{i+1})$. Then,*

$$\bar{U}(w) - U(w) \leq u_{i+1}.$$

Proof. We know that

$$\bar{U}(w) \leq \bar{U}(W_{i+1})$$

and that

$$U(w) \geq U(W_i) = \bar{U}(W_i).$$

Since

$$\bar{U}(W_{i+1}) - \bar{U}(W_i) = u_{i+1},$$

the result follows. \square

Since for any mean-zero random variable Z , $E[U(w + Z)] \leq \bar{U}(w)$, theorem 34 implies that the utility than an agent derives from a gamble is bounded by the size of the item with the highest utility density that she does not already own.

Corollary 5. *Let Z be any mean-zero random variable. Suppose that $w \in (W_i, W_{i+1})$ for some i . Then,*

$$E[U(w + Z)] - U(w) \leq u_{i+1}.$$

Proof. We have

$$E[U(w + Z)] \leq E[\bar{U}(w + Z)] \leq \bar{U}(w)$$

since \bar{U} is concave. \square

Stated another way, the agent always prefers reaching the next high utility density item (i.e. reaching W_{i+1}), to any mean-zero gamble. Compare corollary 5 to corollary 4. These results concern two different ways of measuring the desirability of a gamble: the expected utility and the certainty equivalent. For differentiable utility functions, the derivative provides a link between these two measures—at least for small gambles. However, theorem 32 (and theorem 38 below) tell us that the derivative cannot play the same role when considering knapsack utility functions.

2.5 Approximate concavity at large wealth levels

Our results in this section arise from comparing a knapsack utility function to the utility function arising from the linear relaxation of the underlying knapsack problem. In this section, we will make use of assumption iii: there is some $C > 0$ such that $c_i \leq C$ for all i . Our assumptions on the item set now imply that $u_i \rightarrow 0$, which gives us the following corollaries of theorems 34 and 5. The first tells us that U is very close to \bar{U} at large wealth levels.

Corollary 6. $\lim_{w \rightarrow \infty} \sup_{w' \geq w} |\bar{U}(w') - U(w')| = 0.$

Proof. If $\lim_{i \rightarrow \infty} W_i < \infty$ then the result is trivial since $U(w) = \bar{U}(w)$ if $w > W_i$ for all i . Suppose then that $W_i \rightarrow \infty$. For any $w > W_1$, let

$$i(w) = \sup\{i \in \mathbb{N} : w \geq W_i\}.$$

Then,

$$\begin{aligned} \sup_{w' \geq w} [\bar{U}(w') - U(w')] &\leq \sup_{w' \geq i(w)} [\bar{U}(w') - U(w')] \\ &\leq \sup_{i > i(w)} u_i. \end{aligned}$$

Since the last expression tends to 0 as $w \rightarrow \infty$, the result follows. \square

This result is illustrated by figure 3.2. The second corollary tells us that for large wealth levels, the increase in expected utility (if any) from gambling will tend to be very small.

Corollary 7. *Let Z be a mean-zero random variable. Then,*

$$\limsup_{w \rightarrow \infty} (E[U(w + Z)] - U(w)) \leq 0.$$

Consider the example, illustrated in figure 2.3, in which item i costs $c_i = 10$ and provides utility $u_i = 10/i$. Suppose that the agent has access to the mean-zero gamble Z that pays -5 with probability $1/2$ and pays 5 with probability $1/2$. Consider the case in which the agent has wealth $5 + 10n$ for some $n \in \mathbb{N}$. If she accepts the gamble and loses, her utility will be unchanged since

$$U(5 + 10n) = U(10n).$$

However, if she accepts the gamble and wins, her utility will increase to

$$U(10 + 10n) = U(5 + 10n) + 10/(n + 1).$$

It follows that the agent's increase in expected utility from gambling is

$$E[U(5 + 10n + Z)] - U(5 + 10n) = 5/(n + 1).$$




Figure 2.3: At wealth levels 5, 15, 25, ... the agent will be able to achieve $\bar{U}(5), \bar{U}(15), \bar{U}(25), \dots$ by betting 5 on a fair coin flip. Accepting the gamble increases the agent's expected utility at these wealth levels, but the amount by which the expected utility increases tends to 0 as the wealth level becomes large.

We see this sequence tends to 0, but is always positive. This example shows that it is possible for a wager to provide some incentive to gamble even at very large wealth levels, but corollary 7 ensures that this incentive tends to zero as wealth becomes large.

A slightly different perspective is given by considering our next result, which tells us that the marginal utility of a fixed sum tends to 0 as the initial wealth level becomes large.

Theorem 35. *Let $M > 0$. Then,*

$$\lim_{w \rightarrow \infty} (U(w + M) - U(w)) = 0.$$

Proof. If $\lim_{i \rightarrow \infty} W_i < \infty$, then the result is trivial. Suppose then that $W_i \rightarrow \infty$. Let $w \in [W_i, W_{i+1})$. Then,

$$U(w + M) - U(w) \leq \bar{U}(W_{i+1} + M) - \bar{U}(W_i) \leq d_{i+1}(C + M)$$

and $d_i \rightarrow 0$ as $i \rightarrow \infty$. □

However, the convergence in the theorem need not be monotonic. It is precisely the regions in which we observe a rising marginal utility of wealth that the agent will derive

positive (though possibly very small) incremental utility from gambling, even at large wealth levels.

2.6 Proof that knapsack utility functions are pure jump processes

The proof that a knapsack utility function is a pure jump process will require some preparation. The first step is showing that the supremum in the definition of the knapsack utility function is always achieved.

If x and y are sequences, we will write $x \cdot y = \sum_{i=1}^{\infty} x_i y_i$.

Theorem 36. *For each $w \in [0, \infty)$, there exists an $a \in \{0, 1\}^{\mathbb{N}}$ for which $u \cdot a = U(w)$.*

Proof. Suppose that $(a^n)_{n=1}^{\infty}$ is a sequence of selections for which $c \cdot a^n \leq w$ and $u \cdot a^n \uparrow U(w)$. Since $\{0, 1\}^{\mathbb{N}}$ is compact and metrizable, we may replace $(a^n)_{n=1}^{\infty}$ by a subsequence that converges, say, to a . Fatou's lemma implies that

$$c \cdot a \leq \liminf_{n \rightarrow \infty} c \cdot a^n \leq w,$$

so a is feasible. Similarly, we see that $u \cdot a \leq U(w)$.

Next, let $\epsilon > 0$. Pick K large enough that $w d_K < \epsilon$. Pick N large enough that $a^{n, \leq K} = a^{\leq K}$ for all $n \geq N$. Then, for $n \geq N$, we have

$$\begin{aligned} u \cdot a^{n, > K} &\leq d_K c \cdot a^{n, > K} \\ &\leq d_K \cdot w < \epsilon. \end{aligned}$$

From this we have

$$\begin{aligned} u \cdot a &= u \cdot a^{\leq K} + u \cdot a^{> K} \\ &= u \cdot a^n + u \cdot a^{> K} - u \cdot a^{n, > K} \\ &\geq u \cdot a^n - u \cdot a^{n, > K} \\ &\geq u \cdot a^n - \epsilon. \end{aligned}$$

Letting $n \rightarrow \infty$, we see that $u \cdot a \geq U(w) - \epsilon$ and letting $\epsilon \rightarrow 0$ gives us $u \cdot a \geq U(w)$. \square

A variant of the knapsack problem that we will find useful for the proof is obtained by restricting the item set. Given a list $x = (x_i)_{i=1}^{\infty}$ and a natural number N , we will write $x^{\leq N} = (x_i)_{i \leq N}$. With this notation, we can define the utility of wealth function in the knapsack problem with only the first N items by

$$U^{\leq N}(w) = \sup_{a \in \{0, 1\}^N} \{u^{\leq N} \cdot a : c^{\leq N} \cdot a \leq w\}.$$

We will also be interested in the problem where the item set is restricted to $\{N, N+1, \dots\}$. The notation for this problem is completely analogous to the previous problem.

Our next theorem tells us that for sufficiently large N , $U^{\leq N}$ is a good approximation to U if we restrict our attention to a bounded subset of $[0, \infty)$.

Theorem 37. For any $W > 0$,

$$\sup_{0 \leq w \leq W} |U(w) - U^{\leq N}(w)| \rightarrow 0$$

as $N \rightarrow \infty$. That is, $U^{\leq N}$ converges uniformly to U on compact subsets of $[0, \infty)$.

Proof. Let a be any optimal allocation in the knapsack problem with wealth w . Let a^N be any optimal allocation in the knapsack problem with the item set restricted to the first N items and wealth level w . Then,

$$\begin{aligned} U(w) - U^{\leq N}(w) &= u \cdot a - u \cdot a^N \\ &= (u \cdot a^{\leq N} - u \cdot a^N) + u \cdot a^{>N} \\ &\leq u \cdot a^{>N} \\ &\leq d_{N+1}w \end{aligned}$$

The first inequality follows from the optimality of a^N . The result follows. \square

Proposition. U is càdlàg.⁶

Proof. Theorem 37, tells us that $(U^{\leq N})_{N=1}^{\infty}$ converges uniformly to U on compact subsets of $[0, \infty)$. Each $U^{\leq N}$ is evidently càdlàg, and the uniform limit of càdlàg functions is again càdlàg, so the result follows. \square

Since U is a nondecreasing, càdlàg function, Lebesgue's decomposition theorem tells us that we may write U uniquely as a sum of U_J and U_C where

- (i) U_J and U_C are both nondecreasing and càdlàg;
- (ii) $U_J(0) = U_C(0) = 0$;
- (iii) U_C is continuous;
- (iv) U_J is a pure jump process.

Our next proposition is a reverse dynamic programming equation for solving the knapsack problem given a solution to the tail of the item set.

Proposition. Let N be any non-negative integer. Then, for any wealth level w ,

$$U(w) = \max_{a \in \{0,1\}^N} \{u^{\leq N} \cdot a + U^{>N}(w - c^{\leq N} \cdot a)\}.$$

⁶Càdlàg is an acronym for *continue à droite, limite à gauche*, which is French for right continuous with left limits.

We can now prove theorem 38. The rough idea of the proof is this: Fix $W \geq 0$. We know that $U^{>N}(w) \leq d_{N+1}W$. If there is an interval $[w, w')$ on which the max in proposition 2.6 is achieved by a fixed $a \in \{0, 1\}^N$, then

$$U_C(w') - U_C(w) = U_C^{>N}(w') - U_C^{>N}(w) \leq d_{N+1}(w' - w).$$

If $[0, W)$ can be covered by disjoint intervals for which this is true, then we are able to conclude that

$$U_C(W) \leq d_{N+1}W$$

for all N and so $U_C(W) = 0$.

However, we are not able to guarantee the existence of intervals of this form, so the idea above does not go through. In our proof, we show that it is sufficient to consider intervals $[w, w')$ for which there is some $a \in \{0, 1\}^N$ such that the max in proposition 2.6 is achieved by a at w and at some sequence of wealth levels increasing to w' . We show how to write $[0, W)$ as a disjoint union of intervals in this form.

Proof of theorem 38. Lebesgue's decomposition theorem tells us that it suffices to show that $U_C(W) = 0$ for all $W \geq 0$. Let N be any positive integer and let $W \geq 0$.

Let $w_0 = W$ and let a_0 be an optimal allocation with wealth w_0 . Given w_i and a_i , stop if $w_i = 0$. Otherwise, define $w_{i+1} = c^{\leq N} \cdot a_i$. Let $(w^n)_{n=1}^\infty$ be any sequence such that $w^n \uparrow w_{i+1}$, and for each n let a^n be an optimal allocation with wealth w^n . Define a_{i+1} to be any limit point of $(a^n)_{n=1}^\infty$. Since for each i, j , $a_i^{\leq N} \neq a_j^{\leq N}$ and $\{0, 1\}^N$ is finite, this process must terminate at some stage K . For each $k \in \{0, 1, \dots, K\}$, let

$$U_k(w) = u^{\leq N} \cdot a_k^{\leq N} + U^{>N}(w - c^{\leq N} \cdot a_k^{\leq N}).$$

This is the utility function that arises from imposing the requirement that $a_k^{\leq N}$ be purchased.

For any function f , we will write

$$f(x^-) = \lim_{y \uparrow x} f(y).$$

Our choices of x_k and a_k ensure the following for all $k = 0, 1, \dots, K$:

- (a) $U_k(x_{k+1}) = u^{\leq N} \cdot a_k^{\leq N}$;
- (b) $U(x_k^-) = U_k(x_k^-)$;
- (c) $U_k(x_{k+1}) \leq U(x_{k+1})$;
- (d) $U_k(x_k^-) - U_k(x_{k+1}) = U^{>N}(x_k^- - x_{k+1})$.

We then have

$$\begin{aligned}
U_C(W) &\leq \sum_{k=0}^{K-1} (U(x_k^-) - U(x_{k+1})) && \text{(by neglecting jumps at each } x_k) \\
&\leq \sum_{k=0}^{K-1} (U_k(x_k^-) - U_k(x_{k+1})) && \text{(by (b) and (c))} \\
&= \sum_{k=0}^{K-1} U^{>N}(x_k^- - x_{k+1}) && \text{(by (d))} \\
&\leq \sum_{k=0}^{K-1} U^{>N}(x_k - x_{k+1}) && \text{(since } U^{>N} \text{ is nondecreasing)} \\
&\leq \sum_{k=0}^{K-1} d_{N+1}(x_k - x_{k+1}) \\
&= d_{N+1}W.
\end{aligned}$$

Letting $N \rightarrow \infty$, we see that $U_C(W) = 0$. It follows that $U = U_J$ is a pure jump process. \square

2.7 An intertemporal model with credit constraints

Since Friedman and Savage first explained concurrent gambling and insurance purchase using a utility function with convexities [FS48], several papers have investigated whether or not a convexity in the utility of consumption function in an intertemporal allocation problem results in a convexity in the utility of wealth function in either the discrete or continuous time case. Assuming perfect capital markets, the agent is able to reduce the convexities in the discrete case [BOW80] [HF02] and remove them entirely in the continuous time case [Jon88] by spreading consumption out over time. The convexities that remain in the discrete time case appear to us to be an artifact of discrete time, suggesting that some market imperfection is necessary to elicit convexities in the utility of wealth function.

Any constraint on the agent that makes it more difficult for her to optimally distribute her consumption over time as in [Jon88] is an opportunity for convexities in the utility of consumption function to manifest in the utility of wealth function. We believe that credit constraints represent the most plausible mechanism by which indivisibilities induce convexities in the utility of wealth function. In the presence of credit constraints, the only way for an agent to obtain a costly item is to save for it. If saving is very slow, convexities in the utility of initial wealth function will result.

We will consider a model in which a credit-constrained agent saves over time to purchase items.⁷ We will assume that durable goods retain their full value indefinitely and that the agent faces no transaction costs. Allowing for depreciation or transaction costs would make

⁷We are very grateful to an anonymous referee for suggesting this model.

it more difficult for the agent to spread her consumption over time, resulting in even more pronounced regions of convexity.

Consider a credit-constrained agent with initial wealth w who saves s per unit time, so that the agent's accumulated wealth at time t is $w + st$. The agent may choose from an infinite collection of items as before. Item i costs c_i and provides utility at a rate u_i . The agent discounts future utility exponentially with discount rate r . We will normalize so that item i provides total discounted utility equal to u_i if it is purchased at time 0. Our assumption that durable goods retain their value and that they may be bought and sold costlessly implies that the agent will have wealth $w + st$ at her disposal at time t . As a result, she will allocate her resources so that she achieves utility at the rate $U(w + st)$ at time t . In summary, her utility of wealth at time 0 is given by

$$V_s(w) = \int_0^\infty U(w + st)re^{-rt} dt.$$

The following theorem will help us understand the behavior of V_s .

Theorem 38. *The knapsack utility function U is a pure jump process. That is, there are sequences of point $\{q_j\}_{j=1}^\infty$ and $\{\Delta_j\}_{j=1}^\infty$ such that*

$$U(w) = \sum_{j=1}^\infty \Delta_j 1_{(q_j \leq w)}$$

The number Δ_j is the size of the agent's increase in the rate of utility when she reaches wealth q_j . After some elementary calculus, theorem 38 implies that we may write

$$V_s(w) = \sum_{j=1}^\infty \Delta_j \min\{e^{(w-q_j)/s}, 1\}.$$

We may interpret the function $\min\{e^{(w-q_j)/s}, 1\}$ as a less convex version of the step function that jumps from 0 to 1 at $w = q_j$. In figure 2.4, we plot V_s for various values of s using $r = 0.1$ and the knapsack utility function from figure 2.2. The most important factor in determining the eventual utility that an item provides is the amount of time that the agent must save before purchasing it. When $s = 50$ an additional 50 in initial wealth will avoid only a year of discounting, whereas the same additional 50 would prevent 5 years of discounting when the agent is saving at the rate $s = 10$. If we regard 90% as a meaningful yearly discount rate, we see that significant convexities appear even when the agent wishes to purchase an item that requires only a few years of saving.

The illustration suggests that the agent's utility of wealth function approaches the knapsack utility function as her saving rate becomes very small. This is in fact the case.

Theorem 39. $\lim_{s \rightarrow 0} V_s(w) = U(w)$ for all $w \geq 0$.




Figure 2.4: If the agent values utility a year from now at about 90% of current utility, we see that V_s very close to U when s is small is approximately concave when s is large.

Proof of theorem 39. This follows from our representation of U in theorem 38, our representation of V_s as

$$V_s(w) = \sum_{j=1}^{\infty} \Delta_j \min\{e^{(w-q_j)/s}, 1\},$$

and the fact that $\min\{e^{(w-q_j)/s}, 1\}$ converges to 0 when $w < q_j$ and converges to 1 when $w \geq q_j$. \square

Consider the behavior of an agent with a low saving rate who finds that at some wealth level w , there is a gamble with negative expected value that maximizes her expected utility. If she loses, she will either save or choose to gamble again. If she manages to recoup her losses from the first gamble via saving, she will choose to gamble again when her wealth reaches w . It is also possible that she will choose to gamble again before she has fully recouped her losses. In either case, we expect to see the agent gambling repeatedly until she wins some gamble that increases her wealth above w .

We turn now to the subject of low saving rates. The most straightforward reason for a low saving rate is simply that the agent has a low income. Another reason that may be more plausible for agents with higher wealth levels is the use of hyperbolic discounting. A more sophisticated (and much less tractable) model might allow the agent to choose

between current consumption on non-durable goods and saving for durable goods. As in O'Donoghue and Rabin [OR00], a sophisticated agent who uses hyperbolic discounting may wish (at time 0) to save money for a large purchase, but will correctly anticipate that at some intermediate time she will spend the money on some other purchase. If the agent prefers current consumption to the anticipated consumption at the intermediate time, she will consume and fail to save as a result.

Our results in this section do not depend crucially on the specific saving or discounting scheme that we used. Any wealth accumulation scheme coupled with discounting that has utility from future consumption decreasing to 0 as the consumption becomes more distant will yield the same result as the rate of wealth accumulation decreases. In particular, theorem 39 will continue to be true when the agent uses hyperbolic discounting.

2.8 Conclusion

We introduced knapsack utility functions to serve as a benchmark for the case in which all consumption goods are indivisible. We saw that the convexities that are typical of these functions are more pronounced at low wealth levels. An implication of this is that gambling is a greater source of utility at low wealth levels than at high wealth levels. We then considered the effect of indivisibility size on gambling behavior. We saw that even at large wealth levels, where gambling provides relatively little utility, an agent may very well spend a large sum of money on gambling if she is facing a large indivisibility.

Given that poorer agents derive more expected utility from gambling than wealthy agents, we would expect for poorer agents to gamble at least as often as more wealthy agents. Moreover, we would expect to see a poorer agent expending more effort to gamble, but for the amount gambled to be more correlated with the size of the relevant indivisibilities than the wealth of the gambler.

We hope that our examples have impressed upon the reader the importance of understanding the specifics of the items underlying the agent's consumption decision. The decision about whether or not a particular gamble is worthwhile may change drastically with relatively small changes in the agent's wealth level. Our results indicate that a belief that an agent has a concave utility function when she in fact has a knapsack utility function may lead to wildly incorrect predictions.

We showed that a credit-constrained agent who saves to purchase goods will face significant convexities in her utility of wealth function when her saving rate is low. This effect is in addition to our earlier observation that the large increases in expected utility from gambling will tend to be more common at low wealth levels. Our observation that the agent may find herself again and again in a situation in which participating in a gamble with negative expected value is optimal suggests that indivisibility-induced gambling has the potential to prevent upward mobility. Based on our analysis, we will suggest several targets for removing this barrier.

First, we may attempt to decrease the size of indivisibilities. For example, we might encourage the acceptance by employers of degrees from two-year universities or we might

ensure that low-cost neighborhoods have quality schools. Second, we may attempt to increase saving rates. This could be accomplished by attempting to increase incomes or by increasing the percentage of income saved. Third, we may attempt to make agents more patient. Decreasing the rate at which future utility is discounted will result in less pronounced convexities. If our analysis accurately reflects the economic reality, the practice of filling state coffers with gambling revenue would appear to be a highly regressive practice. We highly recommend [McC94] for a non-technical analysis of some of the practical and ethical considerations of indivisibility-induced gambling.

Chapter 3

Information Aggregation in Financial Markets

3.1 Introduction

Since it was introduced, the Black-Scholes [BS73] framework for modeling security price processes has become ubiquitous. However, this framework is agnostic about the motivations for the decisions made by market participants that drive price changes. Moreover, it is unclear how information about the real economy affects the price of a security.

We introduce a model in which the agents start out believing that the price of the security is governed by a stochastic differential equations as in Black-Scholes. Each agent then observes a private information stream and continuously updates her beliefs about the price process. This spurs each agent to change her trading strategy. The aggregate of the trading strategies then drives the price process. This model allows for events in the real economy to influence prices as long as they show up in the agents' information streams.

In our model, the price process itself is determined by two factors: the beliefs of the agents and the information that the agents receive. The price is not determined solely by beliefs about the fundamentals of a security, but by beliefs about the distribution of the entire price process. Stated another way, if we imagine holding fixed the information streams available to the agents, the security price process will be determined entirely by the beliefs that the agents hold about the security price process. This begs the central question of this chapter: when will agents' beliefs about the distribution of the security price process lead them to behave in a way that makes their initial beliefs correct? That is, when are the agents' beliefs self-fulfilling? We find that the condition that the agents' prior beliefs be correct allows us to draw strong conclusions about the price processes.

Solving each agent's optimization problem after she has observed her private information stream involves representing the price process as an Ito process with respect to the filtration generated both by the market processes and the information stream. Work with the theory of enlargement of filtrations was initiated by Jacod [Jac85], but is able to treat only the case in which the enlargement is achieved by adding a single random variable at time 0.

The work of Ankirchner, Dereich, and Imkeller [ADI07] is sufficiently general to allow the types of enlargements that we are interested in. We primarily make use of the later work of Pikovsky and Karatzas [PK96], in which they show how we may solve optimization problems in continuous time finance after having made an enlargement of filtration.

Our model may be regarded as a model of widespread insider trading. Certainly, we have benefitted from exposure to the Kyle model [Kyl85] and its continuous time analogue [Bac92]. The use of enlargement of filtration in insider trading models is well established [CÇD11, Dan10, FWY99]. Our model differs from these primarily in the sense that it is not game-theoretic. Our agents behave intelligently in the sense that they make difficult optimization decisions in the face of uncertainty, but they do not behave strategically as they do not attempt to respond to the other agents' behavior. Our model has the flavor of a self-confirming equilibrium model as in [FL93], but is perhaps best viewed as a mean-field game [LL07] with a small (i.e., finite) number of players.

The rest of the chapter is organized as follows. In section 3.2 we present our model. In section 3.3, we examine our model in the case that agents use CARA utility functions. In section 3.4, we consider the case of CRRA utility functions. In section 3.5, we conclude.

3.2 The model

There are N agents and a security whose price process is $\{P_t\}_{t \geq 0}$. Each agent also has access to a risk-free bond whose rate of return is normalized to 0. Each agent seeks to maximize her expected utility from wealth at some common end date T .¹ We will make specific assumptions about the utility of wealth functions below.

Suppose that the agents start out with a prior belief under which we may write

$$dP_t = P_t(r_t dt + \sigma dB_t)$$

where $\{B_t\}_{t \geq 0}$ is a Brownian motion with respect to the completed filtration $\mathcal{F} = \{\mathcal{F}_t\}_{t \geq 0}$ that it generates and $\{r_t\}_{t \geq 0}$ is an \mathcal{F} -adapted stochastic process.

Each agent's private information stream is modeled by a complete, right-continuous filtration $\mathcal{G}^n = \{\mathcal{G}_t^n\}_{t \geq 0}$ extending the filtration \mathcal{F} .

Assumption 1. *For each agent n , there are \mathcal{G}^n -adapted processes α^n and \tilde{B}^n such that*

(i) \tilde{B}^n is a \mathcal{G}^n -Brownian motion;

(ii) $dB_t = \alpha_t^n dt + d\tilde{B}_t^n$; and

(iii) $E \left[\int_0^T (\alpha_t^n)^2 dt \right] < \infty$.

¹Setting an end date makes it easy for us to define the problem, but our results are independent of T in the sense that a different choice $T' > T$ for the end date would yield exactly the same processes on the interval $[0, T]$.

Heuristically, this assumption is satisfied when agents retain some uncertainty about the value of the price process at any point in the future even after they observe their private information streams. For details, see [ADI07] or [PK96].

Agent n , working with the filtration \mathcal{G}^n , may write

$$dP_t = P_t(r_t dt + \sigma \alpha_t^n dt + \sigma d\tilde{B}_t^n).$$

We will assume that the agent has access to a risk-free bond with interest rate normalized to 0. The agent believes herself a price-taker, so this is now a classic portfolio optimization problem. The agent's optimal portfolio (the optimal number of shares to hold) at time t will be denoted by θ_t^n . In general, θ_t^n will be a function of the current price P_t of the security. To determine this price, we will enforce the market clearing condition and normalize the number of shares of the security to one:

$$1 = \sum_{n=1}^N \theta_t^n.$$

We will now consider separately the cases in which agents use utility functions with constant absolute risk aversion (CARA) or with constant relative risk aversion (CRRA).

3.3 CARA utility functions

In the CARA case, we will assume that each agent uses the utility function $u^n(c) = 1 - e^{-\gamma c}$ with fixed parameter $\gamma \in \mathbb{R}$ for all agents.

Theorem 40. *Suppose that the agents use CARA utility functions. Suppose also that the agents' prior belief about the distribution of the security price process is correct.*

(i) *The process B is a Brownian motion with respect to the filtration $\mathcal{G} = \bigwedge_{n=1}^N \mathcal{G}^n$.*

(ii) *The drift process r satisfies*

$$dr_t = r_t(r_t dt + \sigma dB_t).$$

(iii) *The drift process $\{r_t\}_{t \geq 0}$ has the explicit solution*

$$r_t = \frac{\exp\{-\sigma^2 t/2 + \sigma B_t\}}{r_0^{-1} - \int_0^t \exp\{-\sigma^2 s/2 + \sigma B_s\} ds}.$$

(iv) *The drift process either converges to 0 or diverges when $r_0 \geq 0$. Moreover, if the map*

$$\phi(x) = P(r_t \rightarrow 0 | r_0 = x)$$

is twice differentiable on $(0, \infty)$ and continuous on $[0, \infty)$, it must take the form

$$\phi(x) = \exp\left(-\frac{2}{\sigma^2}x\right).$$

(v) The price process satisfies

$$P_t = N \frac{r_t}{\gamma \sigma^2}.$$

We will make a few comments before proving the theorem. First, (i) tells us that the process B is unpredictable given only the information that all market participants have in common. However, it is possible that every agent knows something that the others do not, so each may have a nonzero expected drift α_t^n . The forms of (ii), (iv) and (v) are dependent on our choice of CARA utility functions.




Figure 3.1: Sample paths of r with CARA utility functions, $r_0 = 0.05$, and $\sigma = 0.1$.

The conclusion that r_t either converges to 0 or diverges is a bit odd. The belief that the price of a security will not rise faster than the risk free rate will certainly be self-perpetuating provided that the market does not receive any information that contradicts this. The agents are risk-averse, so they will not be willing to tolerate the variance of the security without being compensated, so this means that the price must be tending to zero in tandem with the drift. We see then that (iv) tells us the probability that the security will be worthless from a starting expected rate of return.

Second, the belief that a high drift is a precursor to a yet-higher drift will allay fears that the security may be overpriced, allowing for a perpetual increase in the price. It is unclear to us which property of CARA utility functions is driving this behavior as we fail to observe this kind of perpetual speculative boom in the CRRA case below.

We now present the proof.

Proof. We find that the optimal trading strategy² of agent n , given the filtration \mathcal{G}^n , is given by

$$\theta_t^n = \frac{r_t + \sigma \alpha_t^n}{\gamma \sigma^2 P_t}.$$

If we have some quantity b_t^n for each agent, then we will denote the value of that quantity averaged across agents by \bar{b}_t or by $\overline{b_t^n}$:

$$\bar{b}_t = \overline{b_t^n} = \frac{1}{N} \sum_{n=1}^N b_t^n.$$

Applying our market clearing condition, we see that

$$\begin{aligned} 1 &= \sum_{n=1}^N \theta_t^n \\ &= N \frac{r_t + \sigma \bar{\alpha}_t}{\gamma \sigma^2 P_t}. \end{aligned}$$

Solving for P_t we have

$$P_t = N \frac{r_t + \sigma \bar{\alpha}_t}{\gamma \sigma^2}.$$

Then, we must have that $\bar{\alpha}_t$ is \mathcal{F} -adapted (since we can solve the previous equation to write $\bar{\alpha}_t$ in terms of \mathcal{F} -adapted processes). We then have

$$\begin{aligned} \bar{\alpha}_t &= E[\bar{\alpha}_t | \mathcal{F}_t] \\ &= \frac{1}{N} \sum_{n=1}^N E[\alpha_t^n | \mathcal{F}_t] \\ &= 0. \end{aligned}$$

This implies (v). At this point, the price process is determined by two equations: $\bar{\alpha}_t = 0$ and

$$P_t = N \frac{r_t}{\gamma \sigma^2}.$$

Let us focus on the first of these for a moment. We know that

$$dB_t = \alpha_t^n dt + d\tilde{B}_t^n.$$

Averaging over the agents tells us that

$$dB_t = \bar{\alpha}_t dt + \overline{d\tilde{B}_t^n} = \overline{d\tilde{B}_t^n}.$$

²See [Kar89] for solving general optimization problems in the context of continuous time finance and [PK96] for solving them in the context of enlargements of filtration.

Rewriting this in integral form, we have for $s \leq t$

$$B_t - B_s = \frac{1}{N} \sum_{n=1}^N (\tilde{B}_t^n - \tilde{B}_s^n).$$

Recall that $\mathcal{G}_t = \bigwedge_{n=1}^N \mathcal{G}_t^n$. Since each \tilde{B}^n is a \mathcal{G}^n -Brownian motion, we have

$$E \left[\tilde{B}_t^n - \tilde{B}_s^n \mid \mathcal{G}_s \right] = E \left[E \left[\tilde{B}_t^n - \tilde{B}_s^n \mid \mathcal{G}_s^n \right] \mid \mathcal{G}_s \right] = 0.$$

It follows that

$$E[B_t - B_s \mid \mathcal{G}_s] = 0.$$

Since B_s is \mathcal{G}_s -measurable we obtain (i).

We turn now to the equation

$$P_t = N \frac{r_t}{\gamma \sigma^2}.$$

This implies that

$$dP_t = \frac{N}{\gamma \sigma^2} dr_t.$$

In view of the equation

$$dP_t = P_t(r_t dt + \sigma dB_t)$$

we obtain the equation in (ii):

$$dr_t = r_t(r_t dt + \sigma dB_t).$$

A straightforward, but lengthy, application of Ito's lemma implies that (iii) is in fact a solution to this stochastic differential equation.

Let $\phi(x)$ be the probability that r_t converges to 0 when $r_0 = x$. Since ϕ is assumed to be twice differentiable, we may write

$$\begin{aligned} d\phi(r_t) &= \phi'(r_t) dr_t + \frac{1}{2} \phi''(r_t) (dr_t)^2 \\ &= \phi'(r_t) r_t^2 dt + \phi'(r_t) r_t \sigma dB_t + \frac{1}{2} \phi''(r_t) r_t^2 \sigma^2 dt \\ &= \left(\phi'(r_t) + \frac{\sigma^2}{2} \phi''(r_t) \right) r_t^2 dt + \phi'(r_t) r_t \sigma dB_t. \end{aligned}$$

Since

$$\phi(r_s) = P(\{r_t \rightarrow 0\} \mid \mathcal{F}_s) = E[1_{(r_t \rightarrow 0)} \mid \mathcal{F}_s],$$

we see that $\phi(r_t)$ is a square integrable martingale with respect to the filtration \mathcal{F} generated by B . It follows that the drift term in $d\phi(r_t)$ must be 0 almost surely. Since $r_t > 0$ for all $t \geq 0$ (unless $r_0 = 0$), we see that ϕ must satisfy the differential equation

$$\phi' + \frac{\sigma^2}{2} \phi'' = 0.$$

This is straightforward to solve with the boundary conditions $\phi(0) = 1$ and $\lim_{x \rightarrow \infty} \phi(x) = 0$. The solution is

$$\phi(x) = \exp\left(-\frac{2}{\sigma^2}x\right).$$

This proves (iv). \square

3.4 CRRA utility functions

In this section, we will consider the case in which each agent uses a CRRA utility function. That is,

$$u^n(c) = \frac{c^{1-\gamma} - 1}{1-\gamma}$$

for some $\gamma \in \mathbb{R}$.³ Define X_t^n to be the market value of agent n 's security and bond holdings at time t . We will write

$$X_t = \sum_{n=1}^N X_t^n$$

for the total wealth of the market. We see that X is an \mathcal{F} -adapted process since $dX_t = dP_t$.

To make headway with CRRA utility functions, we will also need to assume that the agent only forms a belief about the drift α_t^n at time t . That is, the agents' private information streams give them information about the instantaneous future, but the agents remain in the dark about anything beyond that.

Theorem 41. *Suppose that the agents use CRRA utility functions, that the agents' prior belief about the distribution of the security price process is correct, that α_t^n is independent of $\mathcal{G}_{t-}^n = \sigma(\bigcup_{s < t} \mathcal{G}_s^n)$. Then,*

(i) *The process B is still a Brownian motion with respect to the filtration $\mathcal{G} = \bigwedge_{n=1}^N \mathcal{G}_t^n$.*

(ii) *The drift process $\{r_t\}_{t \geq 0}$ satisfies*

$$dr_t = r_t \left(1 - \frac{r_t}{\gamma\sigma^2}\right) (r_t dt + \sigma dB_t).$$

(iii) *The drift process $\{r_t\}_{t \geq 0}$ either converges to 0 or to $\gamma\sigma^2$. Moreover, if the function*

$$\phi(x) = P(r_t \rightarrow 0 | r_0 = x)$$

is twice differentiable on $(0, \gamma\sigma^2)$ and continuous on $[0, \gamma\sigma^2]$, then

$$\phi(x) = \left(1 - \frac{x}{\gamma\sigma^2}\right)^{2\gamma+1}$$

for all $0 \leq x \leq \gamma\sigma^2$.

³If $\gamma = 1$, we set $u^n(c) = \log(c)$ as usual.

(iv) The price process satisfies

$$P_t = X_t \frac{r_t}{\gamma \sigma^2}.$$

The result that B is a Brownian motion with respect to the “market filtration” \mathcal{G} is essential for the rest of the results and our proof of this result is dependent on our assumption that agent’s receive only instantaneous glimpses into the future (i.e. that α_t^n is independent of \mathcal{G}_{t-}^n).

In the CRRA case, the long run behavior of the rate of return is much more reasonable than in the CARA case. One might think that the fact that the price remains finite is due to risk aversion. In fact, if the price were to increase rapidly, the agents would become wealthier and, consequently, less risk averse. We would expect for this to lead to *more* risk taking rather than less. A good explanation as to why CARA utility functions result in a speculative bubble whereas CRRA utility functions do not is wanting.




Figure 3.2: Sample paths of r with CRRA utility functions, $r_0 = 0.05$, and $\sigma = 0.1$.

The alternative to a collapse in the price is that the drift tend to $\gamma \sigma^2$. This meshes well with the intuition that γ is the required relative return on an investment to justify a unit of variance. That is, $\gamma \sigma^2 = r_t$. Another interesting feature of this long-run equilibrium is that we will have $P_t \approx X_t$. For this to be the case, we must have almost full investment in the security.

We now present the proof.

Proof. The optimal \mathcal{G} -adapted trading strategy for agent n is⁴

$$\theta_t^n = X_t^n \frac{r_t + \sigma \alpha_t^n}{\gamma \sigma^2 P_t}.$$

In this section, when we have a quantity b_t^n for each agent, we will define

$$\bar{b}_t = \bar{b}_t^n = \frac{1}{X_t} \sum_{n=1}^N X_t^n b_t^n.$$

In the previous section, we averaged over each agent equally. In this section, we are weighting each agent by her wealth.

Now, when we apply the market clearing condition we obtain

$$\begin{aligned} 1 &= \sum_{n=1}^N \theta_t^n \\ &= X_t \frac{r_t + \sigma \bar{\alpha}_t}{\gamma \sigma^2 P_t}. \end{aligned}$$

Solving for P_t we have

$$P_t = X_t \frac{r_t + \sigma \bar{\alpha}_t}{\gamma \sigma^2}.$$

As before, we see that $\bar{\alpha}_t$ is \mathcal{F} -adapted. We then have

$$\begin{aligned} \bar{\alpha}_t &= E[\bar{\alpha}_t | \mathcal{F}_t] \\ &= \frac{1}{X_t} \sum_{n=1}^N E[\alpha_t^n X_t^n | \mathcal{F}_t]. \end{aligned}$$

Now, since \mathcal{F}_t is generated by the continuous process B_t , we have $\mathcal{F}_t = \mathcal{F}_{t^-}$. It follows that $\mathcal{F}_t \subseteq \mathcal{G}_{t^-}^n$. That and the fact that α_t^n is independent of $\mathcal{G}_{t^-}^n$ imply that

$$\begin{aligned} E[\alpha_t^n X_t^n | \mathcal{F}_t] &= E[E[\alpha_t^n X_t^n | \mathcal{G}_{t^-}^n] | \mathcal{F}_t] \\ &= E[X_t^n E[\alpha_t^n | \mathcal{G}_{t^-}^n] | \mathcal{F}_t] \\ &= E[X_t^n E[\alpha_t^n] | \mathcal{F}_t] \\ &= 0, \end{aligned}$$

where the second equality makes use of the fact that X_t^n is continuous and \mathcal{F} -adapted. This implies that $\bar{\alpha}_t = 0$.

Once again, the price process is determined by two equations:

$$\bar{\alpha}_t = 0, \quad P_t = X_t \frac{r_t}{\gamma \sigma^2},$$

⁴For details, see [Kar89] and [PK96].

which gives us (iv).

Taking the wealth-weighted average of the equation

$$dB_t = \alpha_t^n + d\tilde{B}_t^n$$

over all agents yields the equation

$$dB_t = \frac{1}{X_t} \sum_{n=1}^N X_t^n d\tilde{B}_t^n.$$

In integral form this equation is

$$B_t - B_s = \sum_{n=1}^N \int_s^t \frac{X_t^n}{X_t} d\tilde{B}_t^n.$$

Let $\mathcal{G} = \bigwedge_{n=1}^N \mathcal{G}^n$. Then, since \tilde{B}_t^n is a \mathcal{G}^n -martingale and X_t^n/X_t is bounded and \mathcal{G}_t^n -adapted, we have

$$E \left[\int_s^t \frac{X_t^n}{X_t} d\tilde{B}_t^n \middle| \mathcal{G}_s \right] = E \left[E \left[\int_s^t \frac{X_t^n}{X_t} d\tilde{B}_t^n \middle| \mathcal{G}_s^n \right] \middle| \mathcal{G}_s \right] = 0.$$

It follows that

$$E[B_t - B_s | \mathcal{G}_s] = 0.$$

Since B_s is \mathcal{F}_s measurable and $\mathcal{F}_s \subseteq \mathcal{G}_s$, this implies that $E[B_t | \mathcal{G}_s] = B_s$. This proves (i).

We return to the equation:

$$P_t = X_t \frac{r_t}{\gamma \sigma^2}.$$

Writing this equation in differential form, we have

$$dP_t = \frac{1}{\gamma \sigma^2} (X_t dr_t + r_t dX_t).$$

We know that $dX_t = dP_t$, so we may write

$$dP_t = \frac{1}{\gamma \sigma^2} [X_t dr_t + r_t dP_t].$$

We also have

$$\begin{aligned} dP_t &= P_t(r_t dt + \sigma dB_t) \\ &= \frac{X_t r_t}{\gamma \sigma^2} (r_t dt + \sigma dB_t). \end{aligned}$$

Setting these last two equal to each other gives us

$$X_t dr_t + r_t dP_t = X_t r_t (r_t dt + \sigma dB_t).$$

Substituting in the equation for dP_t again gives us

$$X_t dr_t + r_t \frac{X_t r_t}{\gamma \sigma^2} (r_t dt + \sigma dB_t) = X_t r_t (r_t dt + \sigma dB_t)$$

We can divide everything by X_t and rearrange to find

$$\begin{aligned} dr_t &= \left(r_t - r_t \frac{r_t}{\gamma \sigma^2} \right) (r_t dt + \sigma dB_t) \\ &= r_t \left(1 - \frac{r_t}{\gamma \sigma^2} \right) (r_t dt + \sigma dB_t) \end{aligned}$$

This proves (ii).

Let $\phi(x) = P(r_t \rightarrow 0 | r_0 = x)$. Then, using our assumption that ϕ is twice differentiable, we may write

$$\begin{aligned} d\phi(r_t) &= \phi'(r_t) r_t \left(1 - \frac{r_t}{\gamma \sigma^2} \right) (r_t dt + \sigma dB_t) + \frac{1}{2} \phi''(r_t) r_t^2 \left(1 - \frac{r_t}{\gamma \sigma^2} \right)^2 \sigma^2 dt \\ &= r_t^2 \left(1 - \frac{r_t}{\gamma \sigma^2} \right) \left[\phi'(r_t) + \frac{1}{2} \sigma^2 \left(1 - \frac{r_t}{\gamma \sigma^2} \right) \phi''(r_t) \right] dt + \sigma \phi'(r_t) r_t \left(1 - \frac{r_t}{\gamma \sigma^2} \right) dB_t. \end{aligned}$$

As in the previous section, we notice that $\phi(r_t)$ is a martingale with respect to \mathcal{F} . This implies that the drift term in $d\phi(r_t)$ must be 0 almost surely. For $r_t \in (0, \gamma \sigma^2)$, this gives us the differential equation

$$0 = \phi'(x) + \frac{1}{2} \sigma^2 \left(1 - \frac{x}{\gamma \sigma^2} \right) \phi''(x).$$

This equation is separable. We solve to find

$$\phi(x) = \left(1 - \frac{x}{\gamma \sigma^2} \right)^{2\gamma+1}$$

for $0 \leq x \leq \gamma \sigma^2$. For $r_0 = x < 0$, we see that r_t is a bounded above submartingale whose drift is bounded away from 0 when r_t is bounded away from 0, so $r_t \rightarrow 0$. Similarly, when $r_0 = x > \gamma \sigma^2$, r_t is a bounded below supermartingale whose drift is bounded away from 0 when r_t is bounded away from $\gamma \sigma^2$. This proves (iii). \square

3.5 Conclusion

We have considered a continuous time financial model in which agents make trading decisions based on prior beliefs about the distribution of a security price process after they have observed some private information stream. Under specific assumptions on the utility functions used by the agents, we have found explicit descriptions of the dynamics of the model that make the agents' prior belief about the price process correct. In both of the cases that we considered, we observe that the price process is driven by a process which is a

Brownian motion with respect to the filtration containing all of the events that were known to all agents. That is, according to what is known to everyone, the future is unpredictable.

We also observed that a long run equilibrium in which prices collapse to zero is a possibility in both cases. However, in the CARA case we observed the unrealistic scenario in which the drift and price of the security diverge quickly. In the CRRA case, we fail to observe this scenario, but instead observe the possibility of a long run equilibrium in which the drift converges to a rate that offsets the agents risk aversion and leads to approximately exponential growth of the security price along with full investment in the security.

We believe that this model provides a plausible starting point from which we may ask questions about the way in which prices might conform to a given distributional belief when the form of the belief itself affects the distribution. We have focused our discussion on the Black-Scholes model because it is both popular and tractable. It would be interesting to see a thorough treatment of the case of stochastic volatility.

Bibliography

- [AB06] Charalambos D Aliprantis and Kim Border. *Infinite dimensional analysis: a hitchhiker's guide*. Springer Science & Business Media, 2006.
- [ADI07] Stefan Ankirchner, Steffen Dereich, and Peter Imkeller. Enlargement of filtrations and continuous girsanov-type embeddings. In *Séminaire de probabilités XL*, pages 389–410. Springer, 2007.
- [AHKFL87] Sergio Albeverio, Raphael Høegh-Krohn, Jens Erik Fenstad, and Tom Lindstrøm. Nonstandard methods in stochastic analysis and mathematical physics. *Bull. Amer. Math. Soc*, 17:385–389, 1987.
- [AK81] Elie Appelbaum and Eliakim Katz. Market constraints as a rationale for the friedman-savage utility function. *The Journal of Political Economy*, pages 819–825, 1981.
- [And76] Robert M Anderson. A non-standard representation for brownian motion and itô integration. *Israel Journal of Mathematics*, 25(1-2):15–46, 1976.
- [AT75] Phillip Areeda and Donald F Turner. Predatory pricing and related practices under section 2 of the sherman act. *Harvard Law Review*, pages 697–733, 1975.
- [Aum76] Robert J Aumann. Agreeing to disagree. *The annals of statistics*, pages 1236–1239, 1976.
- [Bac92] Kerry Back. Insider trading in continuous time. *Review of financial Studies*, 5(3):387–409, 1992.
- [Bil13] Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.
- [Bla62] David Blackwell. Discrete dynamic programming. *The Annals of Mathematical Statistics*, pages 719–726, 1962.
- [BM13] Paulo Barelli and Idione Meneghel. A note on the equilibrium existence problem in discontinuous games. *Econometrica*, 81(2):813–824, 2013.

- [BOW80] Martin J Bailey, Mancur Olson, and Paul Wonnacott. The marginal utility of income does not increase: Borrowing, lending, and friedman-savage gambles. *The American Economic Review*, 70(3):372–379, 1980.
- [BS73] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *The journal of political economy*, pages 637–654, 1973.
- [Car09] Guilherme Carmona. An existence result for discontinuous games. *Journal of Economic Theory*, 144(3):1333–1340, 2009.
- [CÇD11] Luciano Campi, Umut Çetin, and Albina Danilova. Dynamic markov bridges motivated by models of insider trading. *Stochastic Processes and their Applications*, 121(3):534–567, 2011.
- [Dan10] Albina Danilova. Stock market insider trading in continuous time with imperfect dynamic information. *Stochastics An International Journal of Probability and Stochastic Processes*, 82(1):111–131, 2010.
- [DM86a] Partha Dasgupta and Eric Maskin. The existence of equilibrium in discontinuous economic games, i: Theory. *The Review of economic studies*, 53(1):1–26, 1986.
- [DM86b] Partha Dasgupta and Eric Maskin. The existence of equilibrium in discontinuous economic games, ii: Applications. *The Review of Economic Studies*, 53(1):27–41, 1986.
- [Dob88] Ian M Dobbs. Risk aversion, gambling and the labour-leisure choice. *Scottish Journal of Political Economy*, 35(2):171–175, 1988.
- [Dur10] Rick Durrett. *Probability: theory and examples*. Cambridge university press, 2010.
- [Fan52] Ky Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences of the United States of America*, 38(2):121, 1952.
- [FL86] Drew Fudenberg and David Levine. Limit games and limit equilibria. *Journal of Economic Theory*, 38(2):261–279, 1986.
- [FL93] Drew Fudenberg and David K Levine. Self-confirming equilibrium. *Econometrica: Journal of the Econometric Society*, pages 523–545, 1993.
- [FS48] Milton Friedman and Leonard J Savage. The utility analysis of choices involving risk. *The journal of political economy*, pages 279–304, 1948.

- [FWY99] Hans Föllmer, Ching-Tang Wu, and Marc Yor. Canonical decomposition of linear transformations of two independent brownian motions motivated by models of insider trading. *Stochastic processes and their applications*, 84(1):137–164, 1999.
- [Gli52] Irving L Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [Hak70] Nils H Hakansson. Friedman-savage utility functions consistent with risk aversion. *The Quarterly Journal of Economics*, pages 472–487, 1970.
- [HF02] Roger Hartley and Lisa Farrell. Can expected utility theory explain gambling? *American Economic Review*, pages 613–624, 2002.
- [HH13] Vicky Henderson and David Hobson. Risk aversion, indivisible timing options, and gambling. *Operations Research*, 61(1):126–137, 2013.
- [Jac85] Jean Jacod. Grossissement initial, hypothèse (h?) et théorème de girsanov. In *Grossissements de filtrations: exemples et applications*, pages 15–35. Springer, 1985.
- [Jon88] Larry Jones. A note on the joint occurrence of insurance and gambling. *Kellogg Graduate School of Management, Northwestern University*, 1988.
- [Jor04] Michael I Jordan. Graphical models. *Statistical Science*, pages 140–155, 2004.
- [K⁺41] Shizuo Kakutani et al. *A generalization of Brouwer’s fixed point theorem*. Duke University Press, 1941.
- [Kar89] Ioannis Karatzas. Optimization problems in the theory of continuous trading. *SIAM Journal on Control and Optimization*, 27(6):1221–1259, 1989.
- [KPP] Hans Kellerer, Ulrich Pferschy, and David Pisinger. Knapsack problems. 2004.
- [KRRS94] Jamie Brown Kruse, Stephen Rassenti, Stanley S Reynolds, and Vernon L Smith. Bertrand-edgeworth competition in experimental markets. *Econometrica: Journal of the Econometric Society*, pages 343–371, 1994.
- [KS12] Ioannis Karatzas and Steven Shreve. *Brownian motion and stochastic calculus*, volume 113. Springer Science & Business Media, 2012.
- [KW82] David M Kreps and Robert Wilson. Sequential equilibria. *Econometrica: Journal of the Econometric Society*, pages 863–894, 1982.
- [Kwa65] Ng Yew Kwang. Why do people buy lottery tickets? choices involving risk and the indivisibility of expenditure. *The Journal of Political Economy*, pages 530–535, 1965.

- [Kyl85] Albert S Kyle. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335, 1985.
- [LL07] Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.
- [LS13] Robert Liptser and Albert N Shiryaev. *Statistics of random Processes: I. general Theory*, volume 5. Springer Science & Business Media, 2013.
- [Mar97] Massimo Marinacci. Finitely additive and epsilon nash equilibria. *International Journal of Game Theory*, 26(3):315–333, 1997.
- [McC94] Edward J McCaffrey. Why people play lotteries and why it matters. *Wis. L. Rev.*, page 71, 1994.
- [MMT11] Andrew McLennan, Paulo K Monteiro, and Rabee Tourky. Games with discontinuous payoffs: a strengthening of reny’s existence theorem. *Econometrica*, 79(5):1643–1664, 2011.
- [MN82] Jean-Francois Mertens and Abraham Neyman. Stochastic games have a value. *Proceedings of the National Academy of Sciences of the United States of America*, 79(6):2145, 1982.
- [MS93] A Maitra and W Sudderth. Finitely additive and measurable stochastic games. *International Journal of Game Theory*, 22(3):201–223, 1993.
- [Oks13] Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- [OR94] Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT press, 1994.
- [OR00] Ted O’Donoghue and Matthew Rabin. The economics of immediate gratification. *Journal of Behavioral Decision Making*, 13(2):233, 2000.
- [Pit81] Carolyn Pitchik. Equilibria of a two-person non-zerosum noisy game of timing. *International Journal of Game Theory*, 10(3-4):207–221, 1981.
- [PK96] Igor Pikovsky and Ioannis Karatzas. Anticipative portfolio optimization. *Advances in Applied Probability*, pages 1095–1122, 1996.
- [Pro11] Pavlo Prokopovych. On equilibrium existence in payoff secure games. *Economic Theory*, 48(1):5–16, 2011.
- [Ren99] Philip J Reny. On the existence of pure and mixed strategy nash equilibria in discontinuous games. *Econometrica*, 67(5):1029–1056, 1999.

- [SW57] Maurice Sion and Philip Wolfe. On a game without a value. *Contributions to the theory of games*, 3:299–306, 1957.
- [YH52] Kôsaku Yosida and Edwin Hewitt. Finitely additive measures. *Transactions of the American Mathematical Society*, 72(1):46–66, 1952.