Chapter 2
Ambiguity, Robustness, and Contract Design

Economic actors are often forced to make choices without full knowledge of the consequences. Consider, for example, the decision whether to take an umbrella with you today. The decision would be easy if you knew whether it will rain. But what should your choice be if you are not sure? Your choice would still be easy if you knew the objective probability that it will rain today. But what if even the probability is not known?

Knight (1921) suggested that individuals behave differently when faced with risk (which entails knowledge of objective probabilities of future outcomes, as in the case of coin toss) and uncertainty, when such probabilities are not known. Later, Savage (1954) argued that there is no difference between the two situations. In the case of uncertainty one should simply form subjective believes about possible outcomes and use them to evaluate the expected utility of different choices. However, future experimental work, in particular by Ellsberg (1961) demonstrated that sometimes an individual's choices cannot be rationalized as a result of maximizing a subjective expected utility functional.

In response to the experimental difficulties of subjective expected utility theory researchers have revived the Knightian concept of uncertainty and developed theories about how people react to it. One of the main insights of these theories is that individuals prefer situation of risk to those of uncertainty, the phenomenon known as *ambiguity aversion.*

Concept of ambiguity aversion was formalized by Gilboa and Schmeidler (1989). An ambiguity-averse individual will react by taking into account the worst possible state that can occur, i.e., what will happen if she goes out without an umbrella.

In economics and finance a key issue is how ambiguity aversion will affect contracts the individuals write. In particular, does ambiguity prevent individuals from writing complete contracts and can it sometimes be beneficial by preventing an individual from being cheated in a contract? Will it make decision-makers adopt riskier policies lead to bankruptcies and crises or make them hedge against the

S. Basov, *Social Norms, Bounded Rationality and Optimal Contracts*,
Studies in Economic Theory 30, DOI 10.1007/978-981-10-1041-5_2

worst possible state and stabilize the economy? More generally, how will it affect equilibria of the economy and the very concept of equilibrium?

Related to such questions, there is a new stream of thought in asymmetric information economies based on non-expected utility theory (ambiguity/robustness). When we deal with uncertainty, the choice of expected utility (EU) plays an important role. Even with the same primitives in an economy, if one computes a certain equilibrium concept with different expected utilities (which means a different functional form, as different expected utilities provide alternative functional forms) one will get different results. But then, which formulation of EU is better? Can one compare expected utilities, and on what criteria?

In this chapter I consider the effects ambiguity has on the equivalence of Walrasian equilibria and core allocations to an ambiguous asymmetric information economy; efficiency of the monetary; robustness of the linear contracts in a moral hazard problem; modified Bayesian updating rule for ambiguous beliefs; and the role of ambiguity on the contract incompleteness.

Before describing models built around the concept of ambiguity aversion, let us address a philosophical issue: how ambiguity aversion fits with the subject matter of this book: social norms and bounded rationality?

Suppose an actor has to choose an act $a \in F$ which, together with a state of the world, $s \in S$ will determine a consequence $c = a(s) \in C$.[1] The decision-maker does not know the state if the world and is characterized by some preferences over C. Savage (1954) postulated certain axioms that allow to evaluate acts according a numerical criterion:

$$U(a) = \int u(a(s))\mathrm{d}\mu(s), \tag{2.1}$$

where $u(c)$ is interpreted as utility of the consequence c, and is unique up to a positive affine transformation and μ is a uniquely determined[2] probability measure that captures the decision-maker's beliefs about possible states of the world.

To arrive at representation (2.1) Savage makes the following assumptions (postulates):

P.1: *A preference relation is a transitive and complete binary relation on F.* This postulate ensures that all the acts are comparable. The second axiom, also known as the sure thing principle, requires that the preference between acts depends only on the consequences in states where the payoffs of the two acts being compared are different. This is the key postulate that drives additive separability across the events.

[1]Space S is assumed to be a measure space and acts are assumed to be measurable functions from S into C.

[2]The uniqueness of the probability in Savage's theory is predicated on the convention that the utility function is state independent. Savage axioms guarantee state independence of preferences, but not of the utility and beliefs separately.

P.2: *For all acts, f, f′, that agree on some event E, preferences among them will not change if they are simultaneously altered on E, while still agreeing with each other.*

The third postulate asserts that the decision-maker cares only about the consequences of her actions, not the actions or states of the world per se.

P.3: *The ordinal ranking of consequences is independent of the event and the act that yield them.*

The fourth postulate requires that the betting preferences are independent of the way the bets are made.

P.4: *For all events E and E′ and constant acts x, y, x′, and y′ such that x is preferred to y and x′ is preferred to y′ act that results in x on E and y on its complement is preferred to the act that results in x on E′ and y on its complement if and only if act that results in x′ on E and y′ on its complement is preferred to the act that results in x′ on E′ and y′ on its complement.*

Intuitively, both preferences reflect the idea that event E is more likely than E'. Postulate five rules out trivial preferences. It is crucial in establishing uniqueness of the probability measure in (2.1).

P.5: *There are some constant acts x and y, such that x is strictly preferred to y.*

The sixth asserts that no consequence is either so bad that the decision-maker will not tolerate any act that exposes it even to slightest chance of it happening, no matter how beneficial its other consequences may be, or so good that the decision-maker will not take an act that slightly reduces its probability, even if this act may help to prevent a disaster.

P.6: *For all acts f, g, and h such that f is strictly preferred to g, there is a finite partition*

$$\{E_i \subset S : \bigcup_{i=1}^{n} E_i = S, E_i \cap E_j = \emptyset \text{ for } i \neq j\}$$

of the set of states such that, for all i, f is strictly preferred to act that coincides with h on E_i and with g on its complement and act that coincides with h on E_i and f on its complement is strictly preferred to g.

This postulate implies that there are infinitely many states of the world and that, if there exists a probability measure representing the decision-maker's beliefs, it must be nonatomic.

Finally, the last postulate states

P.7: *If the decision strictly prefers an act to each of the payoffs of another act on a given event, then the former act is strictly preferred to the latter conditionally on the event.*[3]

[3]Technically, one should speak about the non-null events, but this is not important for the current discussion.

> Savage's theorem stipulates that the preference relation satisfies P.1–P.7 if and only if it can be represented by functional (2.1). Probability measure μ in (2.1) summarizes the decision-maker beliefs and erases any distinction between risk and uncertainty.

Savage presented his seven postulates as defining the concept of rational individual. From such point of view any deviation from subjective expected utility theory, in particular behavior that exhibits ambiguity aversion, is a form of bounded rationality.

However, almost from the moment of it inception, the descriptive validity of Savage's model has been criticized. In particular, the sure thing principle, responsible for additive separability of the functional represented preferences across events got under the fire. Ellsberg[4] (1961) demonstrated using simple experiments that individuals display choice patterns that are inconsistent with the existence of beliefs representable by a probability measure. This work leads to creation on non-expected utility models, which often are also built on axiomatic foundations, and require the decision-maker to maximize a certain functional, though a different one from the SEU functional of Savage's theory. If one sees maximizing behavior as a hallmark of rationality, then these theories will still classify as models of rational choice. Therefore, non-expected utility theories, in particular theory of choice by ambiguity-averse decision-makers can be seen as transitional models between models of Savage rationality and models of bounded rationality that completely dispense with maximizing behavior. It is only logical to start this book from considering the consequences of these models.

2.1 A Model of Ambiguity Aversion

Model of ambiguity employed in this book is due to Gilboa and Schmeidler (1989). To understand the motivation behind the model and the way it differs from the SEU framework, consider the following choice problem, known as the *Ellsberg Paradox* (see, Ellsberg 1961). An urn contains thirty red balls, and sixty green and blue balls, in unspecified proportions; subjects are asked to compare (i) a bet on a red draw versus a bet on a green draw, and (ii) a bet on a red or blue draw versus a bet on a green or blue draw. If the subject wins a bet, she receives ten dollars; otherwise, she receives nothing.

The modal preferences in this example are to prefer betting on red to betting in green in (i), but to prefer betting on green or blue to betting red or blue in (ii). One may try to rationalize it follows: betting on red is "safer" than betting on green, because the urn may actually contain zero green balls; on the other hand, betting on green or blue is "safer" than betting on red or blue, because the urn may contain

[4] I describe briefly Ellsberg's experiments in the next section.

zero blue balls, making total number of red or blue balls only thirty, as compared with sixty green or blue balls.

These choices are, however, inconsistent with the SEU framework. Indeed, according to Savage the decision-maker should assign probabilities to a ball been green, blue, or red, i.e., choose three nonnegative numbers,[5] p_g, p_b, and p_r such that

$$p_g + p_b + p_r = 1.$$

Choice in scenario (i) reveals that the decision-maker believes $p_g < p_r$, but then $p_g + p_b < p_r + p_b$, therefore, the decision-maker should prefer betting on red or blue in (ii), contrary to the observed choice. To make sense of this choice assume that rather than having a single belief system the decision-maker allows for multiple beliefs and then computes the most pessimistic expectation with respect to these beliefs. Assume that the allowable beliefs assign probability 1/3 to the drawn ball being green and allow for any pair of nonnegative p_b and p_r as long as the following condition holds

$$p_b + p_r = \frac{2}{3}.$$

Then, when evaluating choices in (i) and considering betting on green one will compute expected value as if $p_g = 0$, while when evaluating choices in (ii) one will compute expected value as if $p_b = 0$. Note that one does not to entertain a view that $p_g = p_b = 0$ simultaneously. One has simply to evaluate expected utility under the worst possible scenario.

This is the essence of the model of ambiguity aversion proposed by Gilboa and Schmeidler (1989). In their model agent distinguishes between situations where probabilities of unknown prospects are known (risk) and situations where they are unknown (uncertainty). The former are characterized by a single probability distribution, while under the second the decision-maker entertains multiple possible subjective beliefs. In Savages theory such a decision-maker will simply have second order probabilistic beliefs about first order probabilistic beliefs and use formula for full probability to collapse them into beliefs about the outcomes. In Gilboa and Schmeidler, on the other hand, the decision-maker will take expectations with respect to the objective beliefs but use min–max criterion with respect to the subjective ones.

Gilboa and Schmeidler provide axiomatic characterization of their criterion. Most axioms are similar to those of Savage. However, they introduce a new axiom that ensures that risk and preferred to uncertainty. To understand it, suppose that we have two urns, each contains 90 balls, and there are 30 red balls in each urn. There are 60 green and 60 blue balls in total, but the decision-maker does not know how they are distributed among urns. We already know that faced with a choice to bid

[5]In the decision-makers beliefs respect objective information then probability of drawing out green ball should be 1/3. The latter is, however, immaterial for our argument.

on red on green ball being drawn from a particular urn a decision-maker will choose to bid on red. No, suppose that urn itself is chosen as a result of a fair coin toss. So, effectively the decision-maker bets either on red or green among one hundred and eighty balls of which 60 are red and 60 are green. Therefore, we eliminated subjective uncertainty and ended up with a situation of pure risk. It is reasonable to assume that now she will be indifferent between betting on red or green and strictly prefer this to bidding on green drawn from any one of the urns. If the coin in the question had not been fair subjective uncertainty would not have been eliminated but would have been reduced. This is the essence of postulate in Gilboa and Schmeidler.

The pessimism of ambiguity-averse decision-makers might under certain circumstances turn out to be quite useful. For example, it can help to mitigate, and sometimes completely resolve, tensions between efficiency and incentive compatibility, both in the framework of Walrasian equilibria and under bilateral bargaining under private information. We will analyze scenarios where it happens below. However, it can also prevent decision-maker pursuing uncertain but ultimately beneficial courses of action, for example, investing in fundamental research. Therefore, ambiguity aversion shares with other forms of deviations from behavior based on Savage style rationality and individualistic preferences, such as bounded rationality and sensitivity to social norms, the distinction of been both a blessing and a curse. They are blessings since, if properly used, they can mitigate some traditional trade-offs, but a curse since they come with costs of their own. I will return to this topic in Chap. 6, where I discuss interaction of social norms with optimal incentives provision.

2.2 Equilibrium Theory and Ambiguity

This section is based on the work of He and Yannelis (2015) and considers how concepts of core, Walrasian equilibrium, and relation between them can be generalized to an ambiguous asymmetric information economy.

They start by remarking that the standard Arrow-Debreu state contingent model allows the state of nature of the world to be involved in the initial endowments and payoff of the agents, who make contacts ex-ante, i.e., before the state of nature is realized and once the state is realized the contract is executed and consumption takes place. Radner (1968, 1982) extended the analysis of Arrow and Debreu by introducing asymmetric private information. The private information is modeled as a partition of a finite state space, with the requirement that the initial endowments of each agent measurable with respect to sigma-algebra generated by her partition. The issue of incentive compatibility does not arise in this model, as all the contracts signed ex-ante are assumed to be binging. However, for this to make sense one must assume that there is an exogenous court or government that enforces the contract ex-post, otherwise agents may find it beneficial to renege.

Holmström and Myerson (1983) argued that if one assumes that the agents are Bayesian expected utility maximizers, it is not possible to have allocations which are both first best efficient and also ex-post incentive compatible. He and Yannelis asked whether one can find another decision theoretic framework that will allow defining concepts of Walrasian equilibrium and core in the symmetric information economy in such a way that both core and Walrasian equilibria exist, and are both incentive compatible and first best efficient. Their main finding is that such a framework is provided by maximin expected utility, i.e., by exactly the kind of preferences used by Gilboa and Schmeidler to model ambiguity aversion.

To give the reader the main idea of the argument, let us consider the following example, found in He and Yannelis. Consider an economy with one commodity, two agents, indexed 1 and 2 and three states of the world, indexed *a, b,* and *c*. The initial endowments of the agents are

$$e_1 = (5, 5, 0); e_2 = (5, 0, 5).$$

Their information partitions of are given by[6]

$$\Pi_1 = \{\{a, b\}, c\}; \Pi_2 = \{\{a, c\}, b\},$$

i.e., each agent can be either poor or rich. If the agent is poor she knows that the other agent is rich, but the rich agent does not know whether the other agent is rich or poor.

Recall, that Yannelis and He define a tuple consisting of an allocation vector and a vector of prices to be a Walrasian expectations equilibrium (WEE) for the economy, if each agent's allocation is measurable with respect to her information partition and maximizes her expected utility given prices and her beliefs; and the price vector is such the markets clear. Measurability requirement is crucial in this definition. In the example above it can be interpreted in the following way: each agent has to pay a tax/receive a subsidy depending on whether she is rich or poor, but independent on the wealth of the other agent. Now friction between ex-post incentive compatibility and efficiency can be understood and an incentive to lie about once wealth.

Maximin expectation equilibrium (MEE) differs from WEE by assuming that the agents maximize the expected utility under most pessimistic expectations consistent with their signal and the measurability requirement is not imposed. We will see when discussing the example why dropping measurability requirement is reasonable in the case of MEE.

To continue with the example, let us assume that both agents have identical Bernoulli utility function

$$u(x) = \sqrt{x}$$

[6]Note that the endowments are measurable with respect to their partitions.

and the prior beliefs of both agents are the same

$$\mu(\omega) = \frac{1}{3} \text{for } \omega \in \{a, b, c\}.$$

Suppose that agents are both Bayesian expected utility maximizers. It can be easily checked that there is no nonfree disposal Walrasian expectation equilibria with positive prices.

If we allow for free disposal, then

$$x_1 = (4, 4, 1) \text{ and } x_2 = (4, 1, 4)$$

is a (free disposal) WEE allocation with the equilibrium price $p(a) = 0$ and $p(b) = p(c) = \frac{1}{2}$. However, this allocation is not incentive compatible.

On the other hand, if the agents are maximin expected utility maximizers, then there exists maximin expectations equilibrium (y, p), where

$$y_1 = (5, 4, 1), y_2 = (5, 1, 4)$$

and price is given by $p(a) = 0, p(b) = p(c) = \frac{1}{2}$. If state b or c is realized, the ex-post utility of agent one will be the same in both Bayesian preference setting and maximin preference setting, since $x_1(b) = y_1(b)$ and $x_1(c) = y_1(c)$. But if state a occurs, the ex-post utility of agent one with maximin preference will be strictly higher than that in the Bayesian preference setting, since

$$x_1(a) = 4 < 5 = y_1(a).$$

Therefore, the maximin preference allows agents to reach higher efficiency.

To understand the difference between Bayesian and maximin preferences it will be desirable to recast the above example as a social choice problem. There are three states of the world, in one both agents are rich and in the other two one agent is rich and the other is power. A priori the agents will find it beneficial to agree to tax the rich agent and subsidize the poor one. Suppose once endowments are realized the agents must declare whether they are rich or poor. If an agent declares to be rich she has to pay \$1 and then the tax revenue is distributed in the following way: first the agent who declared herself poor receives \$1, and if there are more tax proceeds collected the rich agents also receive \$1. The announcement triggers an audit if and only if the reported state is inconsistent with prior information, i.e., both agents claim to be poor. In this case the social planner finds the truth, and she either taxes the rich agent and pays the poor one or does nothing if both turn to be rich. I would like to argue that truth telling is not Bayes-Nash equilibrium (BNE) under Bayesian preferences, but it is a BNE under the maximin preferences.

Note that the poor agent never has incentive to lie, so she will always announce to be poor. Let us consider incentives of the rich agent. If she claims to be rich, her payoff will be \$5 if the other agent claims to be rich and \$4 if the other agent claims to be poor. If, on the other hand, she claims to be poor, her payoff will be \$4 if the other agent claims to be poor and is found to be poor, \$5 if the other agent claims to be poor, but is found to be rich, and \$6 if the other agent claims to be rich. Since there is positive probability (50 %) that the other agent is in fact reach, a rich agent with Bayesian preferences will strictly prefer to say she is poor. For the agent with maximin preferences, on the other hand, the only contingency that matters is the one when the other agent is poor, therefore, she will be indifferent between telling the truth and lying.

Note, however, that since the lottery generated by lying first order stochastically dominates the one generated by telling the truth, truth telling remains a weakly dominated strategy. One can avoid it by imposing fine if one is audited and caught lying. Then the result can be interpreted as stating that with maximin preferences even a small fine will induce truth telling, since as long as it is possible to be caught such agents will act as if they are certain to be caught. This is the key insight that explains why ambiguity-averse preferences mitigate tension between efficiency and incentive compatibility. It turns out, as I will explain later in this book, that the maximin preferences are the only preferences that eliminate the trade-off completely. However, if preferences are convex combination of maximin and Bayesian, the information rents necessary for truthful revelation will be smaller than for purely Bayesian preferences. He and Yannelis proceed to argue that under the maximin preferences the MEE coincide with private core[7] for large economies.

In conclusion, He and Yannelis developed a new asymmetric information economy framework, which allows for ambiguity-averse preferences. They derived new existence and equivalence results for MEE and private core. The most important insight of this work is that ambiguity aversion eases tensions between efficiency and incentive compatibility. We will see other instances of this insight later in this chapter.

2.3 Ambiguity Aversion and the Myerson-Satterthwaite Theorem

This section is based on De Castro and Yannelis (2011) and discusses the way ambiguity aversion allows one to avoid conclusions of Myerson-Satterthwaite theorem. The conflict between efficiency and incentive compatibility arises in many areas of economics. In particular, it features in auction theory, bargaining theory, and the theory of general equilibrium with asymmetric information, among other cases. One of the key insights we got in the last section is that, at least in the case of general equilibrium models, this conflict can be mitigated if the agents are ambiguity-averse. Moreover, when individuals have maximin expected utility

[7]Definition of private core is similar to that of core, but with crucial requirement that allocations are measurable with respect to private information partitions.

(MEU) preferences, then the conflict can be resolved completely, any efficient allocation is incentive compatible. Conversely, only MEU preferences have this property, though tension is substantially mitigated for much broader class of preferences.

To fix ideas let us start with the following example from De Castro and Yannelis (2011). Consider a seller and a buyer, who both have private valuations of an object. The seller's valuation is $v \in [0;1]$ and the buyer's valuation is $t \in [0;1]$. Trade will result in an efficient allocation if it happens if and only if $t > v$. Assuming both the buyer and the seller are expected utility maximizer's, Myerson and Satterthwaite (1983) have proved that there is no incentive compatible, individual rational, budget balanced mechanism that would achieve ex-post efficiency in this situation.

Let us consider the following mechanism, known as the double auction. The seller places an asking price a and the buyer submits a bid b. If the bid is above the ask, they trade at price

$$p = \frac{a+b}{2}.$$

If the bid is below the asking price, there is no trade. Therefore, if they if the trade occurred at price p, the (ex-post) profit for the seller will be $p - v$, and $t - p$ for the buyer; otherwise both get zero. Since the announcement of the agent affects price at which trade occurs the buyer will have incentives to shed her value and the seller an incentive to exaggerate the cost. The logic behind this result is the same as behind shedding the value in the first price sealed-bid auctions.[8] This means that if the value of the buyer is only slightly above that of the seller the trade will not occur, despite the efficiency gains it will bestow. Myerson and Satterthwaite (1983) argued that under the Bayesian paradigm this efficiency cannot be avoided by using a more sophisticate mechanism as long as it is required to be incentive compatible, individual rational, budget balanced.

One can avoid this conclusion if one dispenses with the assumption of probabilistic sophistication that requires that each agent forms a single prior about the distribution of the other agents' values and allow her to entertain a variety of possible beliefs. This implies that both the buyer and the seller have to make choice in a situation of Knightian uncertainty. Following Gilboa and Schmeidler, we will model it using the maximin criterion.[9]

The maximin criterion implies that each individual considers the worst-case scenario for each action, and chooses the action that leads to the best worst-case outcome. Let us argue that true revelation is incentive compatible under this criterion. Indeed, truthful announcements of $a = v$ by the seller and $b = t$ by the buyer are incentive compatible if buyer and seller do not have any incentive to choose a

[8]This logic assumes that both agents are expected utility maximizers.

[9]Use of the maximin criterion in classical statistics dates back to Wald (1950), but the behaviorial foundations where first provided by Gilboa and Schmeidler (1982).

different action. If the buyer chooses $b = t$, the worst-case scenario is to end up with zero (either by buying by $p = t$ or by not trading). If she chooses $b > t$, the worst-case scenario is to buy by $p > t$, which leads to a (strict) loss. If she considers $b < t$, the worst-case scenario is to get zero (it always possible that there is no trade). Therefore, neither $b < t$ nor $b > t$ is better (by the maximin criterion) than $b = t$ and she has no incentive to deviate. The argument for the seller is analogous.

Note, however, that truthful announcement is as good as any announcement below true value for the buyer and above the cost for the seller. One may argue that telling lies may require more cognitive resources that telling the truth. This will make truth telling strictly dominant under the maximin criterion. Therefore, combination of ambiguity aversion with bounded rationality can more convincingly resolve friction between efficiency and incentive compatibility, then any of them can achieve of their own.

De Castro and Yannelis (2011) also note that another interesting property of these preferences is that the set of efficient allocations is not small. At least in the case of one-good economies, the set of efficient allocations under maximin preferences includes all allocations that are incentive compatible and efficient for the expected utility maximizers. They also argued that under some reasonable conditions imposed on preferences[10] maximin preferences are the only preferences that resolve the conflict between efficiency and incentive compatibility. However, if one settles for mitigation rather than the complete resolution of the conflict, more general ambiguity-averse preferences can help.

Note that ambiguity does not always improve social welfare. Mukerji (1998), for example, argued that in hidden action models[11] it may decrease efficiency, by limiting trading opportunities. We will review that paper in the next section. It may also lead to contractual incompleteness, as noted by Mukerji (1998) and Grant et al. It is also important to mention that the maximin preferences are the only ones that allow to completely resolving the conflict between incentive compatibility and efficiency, though any ambiguity-averse preferences mitigate this conflict.

Other ways to get around the problem of conflict between incentive compatibility and efficiency were proposed in the literature. Yannelis (1991), for example, imposes the private information measurability condition and argues in a series of papers[12] that it forces incentive compatibility of any Pareto optimal allocation. Indeed, if an agent trades a nonmeasurable contract, she effectively makes promises depending on conditions that she cannot verify. Therefore, other agents may have an incentive to cheat her, which leads to the failure of incentive compatibility. Insistence on only measurable contracts preserves incentive compatibility; however

[10]The conditions are rationality, monotonicity, and continuity.

[11]As we will see later, hidden action and hidden information models often respond in different ways to deviations from Bayesian rationality. For example, the costs of boundedly rational behavior are usually borne by the principal in hidden action models, but can be borne by either the principal, or the agent, or both in hidden information models.

[12]See papers by Krasa and Yannelis (1994), Koutsougeras and Yannelis (1993), Hahn and Yannelis (1997).

it restricts trade and may lead even lead to no trade. In financial markets this requirement will mean that the traders cannot use the price of the assets to deduce information possessed by other traders, i.e., it will exclude the procedure which underlies the rational expectation equilibria and leads to the efficient market hypothesis.

Gul and Postlewaite (1992), McLean and Postlewaite (2002) proposed yet different solutions to the conflict. They assumed that the agents are “informationally small” and showed the existence of incentive compatible and approximately Pareto optimal in a replica economy.

The above approaches, however, preserve Bayesian rationality paradigm. Given growing empirical and experimental evidence that calls this paradigm into question, it is important to ask how common and how severe the conflict is. The answer that emerges from this research is: quite common, but probably not as severe as we originally thought.

Finally, it would be interesting to study an evolutionary model of populations of agents with different attitudes to ambiguity. Such a model forms into a more general class of the models of evolution of preferences. Robson provided a comprehensive review of such models. Bounded rationality plays important role in this approach, since fully rational agents will simply have their utility equal to their expected inclusive fitness. Therefore, one should attempt to build a model, where preferences and decision-making rules coevolve. In the next chapter I will make some remarks about how such models can be built.

2.4 Ambiguity Aversion, Moral Hazard, and Contractual Incompleteness

In the previous section we saw examples of situations where ambiguity aversion allows improving efficiency by mitigating tensions between incentive compatibility and efficiency in Walrasian and Myerson-Satterthwaite settings. The common feature of examples considered so far was that economic agents possessed hidden information but did not have to undertake hidden actions. In this section we will consider the situation, when the agents have to undertake such an action, in particular, a relation-specific investment, which will affect the value of a widget for both parties. Building on the ideas developed by Mukerji (1998), Grant et al. (2006), I will argue that in this case ambiguity-averse preferences will lead to incomplete contracts and therefore, loss of efficiency.

Let us start with an example that appears in Mukerji (1998). Consider two vertically related risk-neutral firms, B and S. Assume that the set of possible states of the world, Ω, contains three elements, i.e.,

$$\Omega = \{\omega_0, \omega_b, \omega_s\}.$$

At date zero each firm decides on the level of relation-specific investment, which can be either high or low. Let $\beta_L(\sigma_L)$ and $\beta_H(\sigma_H)$ denote low and high levels of investment for firm $B(S)$, respectively. The social surplus in the states of the world is given by

$$s(\omega_0) = 0, s(\omega_b) = s(\omega_s) = 200$$

and the costs of investment are

$$h_B(\beta_L) = h_S(\sigma_L) = 10, h_B(\beta_H) = h_S(\sigma_H) = 85.$$

As before, we assume that agents maximize minimal expected utility with respect to a set of beliefs. To formalize the idea of allowable beliefs Mukerji makes use of idea on nonadditive measure (capacity), which is defined in the following way.

Definition Let Ω be a finite set. A function $\pi : 2^\Omega \rightarrow [0, 1]$ is called a nonadditive measure (capacity) if it has the following properties: (i) $\pi(\emptyset) = 0$, (ii) $\pi(\Omega) = 1$, (iii) for any $A, B \in 2^\Omega$, $(A \subset B) \Rightarrow (\pi(A) \leq \pi(B))$. Capacity $\pi(\cdot)$ is called convex if for all $A, B \in 2^\Omega$

$$\pi(A \cup B) \geq \pi(A) + \pi(B) + \pi(A \cap B).$$

One can interpret capacity of a set as the minimal possible probability the decision-maker assigns to the set. Given a convex capacity, the set of possible probabilistic beliefs a decision-maker entertains is given by its core

$$\Pi(\pi) = \left\{\pi_j \in \Delta(\Omega) | \pi_j(X) \geq \pi(X) \text{ for all } X \in 2^\Omega\right\}.$$

As in the previous two sections, the decision-maker makes her choices on the basis of the maximin expected utility, however, the beliefs are restricted to belong to a core of a given capacity, rather than being arbitrary.

Let us, following Mukerji, assume that

$$\pi(\beta_L, \sigma_L) = (0.78,\ 0.01,\ 0.01),$$

i.e., if both buyer and seller choose low levels of investment, state ω_0 will realize with at least 78 % chance, while states ω_b, ω_s will realize with at least 1 % probability each. Similarly,

$$\pi(\beta_H, \sigma_H) = (0.02,\ 0.39,\ 0.39)$$
$$\pi(\beta_H, \sigma_L) = (0.42,\ 0.365,\ 0.015)$$
$$\pi(\beta_L, \sigma_H) = (0.42,\ 0.015,\ 0.365).$$

Therefore, the buyer's effort shifts likelihood from the low surplus state predominately to the state favored by the buyer, while the seller's effort shifts it predominately to the state favored by the seller. Let us further assume that

$$\pi(\{\omega_b, \omega_s\}|\beta, \sigma) - \pi(\{\omega_b\}|\beta, \sigma) - \pi(\{\omega_s\}|\beta, \sigma) = 0.1$$
$$\pi(\{\omega_0, \omega_s\}|\beta, \sigma) - \pi(\{\omega_0\}|\beta, \sigma) - \pi(\{\omega_s\}|\beta, \sigma) = 0$$
$$\pi(\{\omega_b, \omega_0\}|\beta, \sigma) - \pi(\{\omega_b\}|\beta, \sigma) - \pi(\{\omega_0\}|\beta, \sigma) = 0.$$

Also each agent assigns higher minimal probability that a good state will occur than the sum of minimal probabilities that the buyer's and the seller's preferred states will occur. Intuitively, this implies that though each agent is optimistic about the possibility of her investment to result in a high surplus state, both agents are worried that the other guy will benefit.

One can verify by a direct calculation using maximin expected utility criterion that (β_H, σ_H) is the first best, and (β_L, σ_L) is the second-best action profile. Given that each agent's effort mainly increases the likelihood of her preferred high surplus state, it would be natural to guess that the first best can be by allocating the entire surplus at ω_b to the buyer and at ω_s to the seller. Such a contract will, indeed, be incentive compatible. However, no ex-ante transfers may be arranged to make it individually rational. Also, since the incentive constraints bind, any contract which attempts to smoothed the ex-post payoffs would break at least one of the incentive constraints. Therefore, the first best cannot be implemented.

Let us assume that in the case of lack of a contract, the social surplus in the resulting state is split equally. Then it is straightforward to verify that the null contract implements the second-best investment profile. One can see the null contract as a complete contract that instructs that no trade should occur in every contingency, but allows for ex-post renegotiation. Obviously, it is as good as leaving the contingencies unmentioned, i.e., an ultimate incomplete contract. One needs to assume transaction costs (which could be infinitesimally small), to make the incomplete contract strictly optimal.

Another way in which ambiguity aversion could interact with bounded rationality to lead to contractual incompleteness was explored by Grant et al. (2006). Their central idea is that boundedly rational individuals do not have access to a language sufficiently rich to describe all possible states of nature, which leads to ambiguity of a contract. As a result, risk-averse agents may forgo potential gains from risk sharing and choose incomplete contracts instead. It can be illustrated by a simple example.

Suppose two farmers Robin and Clarke are considering a possibility of entering a risk sharing contract. They grow different crops and know that if the weather is rainy then Robin will have a good harvest, and Clarke will get none, and if there it is sunny their fortunes will be reversed. However, each observes only weather on their own farm and is unaware of the possibility that weather may differ at different locations. Been risk averse, they would prefer a priori to share harvest equally, however, being boundedly rational they can only think of a contract of a form: *if it is rainy Robin delivers half of her harvest to Clarke, and if it is sunny Clarke delivers half of her harvest to Robin*.

In the formal framework developed below, if such a contract were signed, the presumption is that each party translates the contingencies on which the transfer function depends into her or his own experience. Therefore, if it is rainy at Robin's farm and sunny at Clarke's or vice versa, the terms of contract will lead to disagreement. The first scenario is less problematic, since both parties are expected to deliver, and therefore dispute can be easily resolved by exchange. In the second case, however, this can lead to a serious dispute. The authors assume that there are costs to dispute resolution.

This example demonstrates that boundedly rational players may be unable to formulate a sufficiently refined description of the states of the world to avoid dispute. However, as argued by Grant et al. (2006), they may be aware that disputes are possible. This may lead them to choose a null contract.

This corresponds closely to the distinction between risk and uncertainty I discussed above. The larger is the gray area giving rise to dispute, the less will parties benefit from a complete contract. On the other hand, they will benefit more from such a contract the more risk-averse. Thus risk and ambiguity work in opposite directions.

In both Mukerji and Grant et al. approaches some transaction costs are necessary to justify contractual incompleteness. However, though transaction costs (which are ultimately a form of bounded rationality) are ultimately necessary to justify contractual incompleteness, ambiguity aversion significantly reduces marginal gains from including more details in the contract and allows one to get away with small (sometimes even infinitesimal) transaction costs. Segal (1999), Hart and Moore (1999) have shown that complexity of the environment can have similar effects. I will discuss that paper in Chap. 4.

2.5 Some Other Economic Effects of Pessimism

Ambiguity aversion can be considered as a form of pessimism concerning unknown probability distributions. This view of ambiguity aversion is supported by psychological studies. Pulford (2009), for example, studied the influences of optimism and pessimism on ambiguity aversion in a standard Ellsberg urn experiment and found that highly optimistic people showed significantly less ambiguity aversion than their pessimistic counterparts, when information was given that the number of balls was randomly determined. When ambiguity is clear, and trust issues are removed, subjects' optimistic outlook influences their degree of ambiguity aversion and thus their decisions. This pattern was present but less pronounced in the condition when the composition of the ambiguous urn could be interpreted as being influenced (rigged) by the experimenter. Pulford has also observed that the perception of the situation, especially the degree of trust in the experimenter, was significantly influenced by the participants optimism. This observation potentially opens the door for modeling of coevolution of the social norms, such as trust, attitudes to ambiguity. I will briefly touch on building of coevolutionary models of preferences, decision rules, and social norms in Chap. 3.

Ambiguity aversion can be considered as a form of pessimism concerning unknown probability distributions. We have seen that consequences of this can be both: positive, for example, allowing for mitigation of the trade-off between incentive compatibility and efficiency, and negative, for example, leading to the contractual incompleteness. Here I will briefly describe some other economic phenomena, which are governed by some form of pessimism of economic actors, though not necessarily by ambiguity aversion.

2.5.1 Robustness and Linear Contracts: Uncertainty Over Agent's Actions

Linear contract are prevalent in economic life. They also are used widely in contract theory and will make use of them later in this book. But what is the ultimate rationale for linearity? Carroll (2015) tackles this question and argues that it may be the principal's uncertainty concerning the set of actions the agent can take. The framework is the following.

A principal hires an agent to perform a costly action on her behalf. Both the principal and the agent are risk neutral. Action is not observable, but gives rise via a stochastic technology to an observable output; $y \in Y$. It is assumed that the set of possible outputs, Y, is a compact subset of the real line with the minimal element normalized to be zero. The agent's action is a pair $(F, c) \in \Im \subset \Delta(Y) \times R_+$, i.e., the agent selects a distribution of outputs at some nonnegative cost, where the set of allowable pairs are given by technology $\Im$, where $\Im$ is assumed to be compact. When choosing an action, the agent optimally responds to the incentive scheme, provided by the principal, where the scheme specifies payment $w(y)$ for the output value y, where $w(\cdot)$ is assumed to be continuous and nonnegative.[13] Let $A^* \subset \Im$ be the set of agent's optimal choices, which is guaranteed to be nonempty.

[13]The latter is the limited liability constraint, which prevents the principal from selling the enterprise to the agent.

If the principal knows the technology then this is a standard principal agent problem. Carroll, however, assumes that the principal does not know the technology and maximizes her minimal payoff, where minimum is taken over all technologies which allow agent to exert no effort and produce no output. Under these assumptions Carroll shows that the optimal contract is the linear one.

Intuition for Carroll's result is easy to understand. Suppose optimal contract (whatever form it might take) achieves expected payoff π_P for the principal and π_A for the agent. This means that the total expected output is given by

$$y = \pi_P + \pi_A.$$

Note that principal can achieve the same outcome by offering the agent linear contract

$$w(y) = \frac{\pi_A}{\pi_P + \pi_A} y.$$

Indeed, since agent had chosen before action that resulted in expected payoff π_A, she will know choose action that results in at least the same expected payoff, and since under the linear scheme payoffs of the principal and the agent are proportional, the principal will end up earning at least π_P. The rigorous argument has to deal with some technical subtleties; interested reader is referred to Carroll (2015) for the details.

Carroll also points out that to explain prevalence of linear contracts in practice via this model one need not interpret it literally. Instead, one may assume that decision-makers are just looking for the simplest contact that guarantees to perform reasonably, akin to Simon's (1956) notion of satisfying, providing further link between models build around the concept of pessimism and models of boundedly rational decision-making. Finally, Carroll notes that this model cannot be used to justify the common practice in applied theory to assume full knowledge of the environment, but assume linearity for tractability, since the optimal linear contract in that case is different from the maximin optimal contract.

2.5.2 *Monetary Equilibria with Wary Agents*

Yet another form pessimism can take is wariness about the future. Araujo et al. (2014) assumed that when choosing their life-time consumption profiles agents neglect gains at distant dates, but take into account losses. They argued that to implement the efficient allocation among wary agents one will require a nonvanishing supply of money.

To model wariness, Araujo et al. assume that utility of an agent is given by

$$U(x) = \sum_{t=1}^{+\infty} \delta^t u(x_t) + \beta \inf_{t \geq 1} u(x_t).$$

Here $\{x_t\}_{t=1}^{\infty}$ is a bounded consumption profile that is financed by the period endowment and fiat money holdings net of tax obligations. Fiat money is injected into the economy at time zero. At the later dates money supply evolves endogenously, with aggregate money supply at a particular date equal to the aggregate money supply at the previous date net of tax payments. Walrasian equilibrium of the economy is defined in the usual fashion.

It is the last term in the utility function, which depends on the potential unfavorable shock in the future that makes fiat money holding attractable. Note that if $\beta = 0$ the agents are conventional discounted expected utility maximizers, and fiat money will not have any value at equilibrium. Intuitively, if one assumes that $\beta > 0$ fiat money holding will allow the consumers the marginal benefit of raising infimum of consumption, provided it is never reached in finite time (the worst is always yet to come). If this cost outweighs the opportunity cost of carrying fiat money, positive money holding become optimal.

2.6 Concluding Remarks

In this chapter I discussed effects of ambiguity aversion and other forms of pessimistic biases on economic decision-making. Discussion in this chapter shows that such behavior can help to mitigate some important trade-offs, such as the trade-off between implementability and Pareto optimality and as a result improve economic outcomes. It can, however, also lead to contractual incompleteness, which in turn leads to hold-up problems and decreases relation-specific investments below the efficient levels. This kind of behavior can also help to explain some common economic phenomena such as simplicity of real life contracts in comparison with the ones suggested by the optimal contract literature and the prevalence of the fiat money.

In the later chapters I will discuss other, more drastic, forms of deviations from the Savage's paradigm and study how they modify the nature of the optimal contracts. In particular, I will note some similarities between the ways complexity and ambiguity aversion affect the structure of the optimal contracts.

Before discussing in depth effects of bounded rationality and social forces on the optimal contracts, I will pause and ask is there a common threat that unites different deviations from the Savage's paradigm. I will investigate this question from the evolutionary point of view and suggest a possible way to model coevolution of bounded rationality and ambiguity aversion.

References

Araujo, R., Gama-Torres, J. P., Novinski, R., Pascoa, M. R. (2014). *On efficiency of monetary equilibria when agents are wary*. http://www.laef.ucsb.edu/pages/conferences/are15/papers/araujo.pdf

Carroll, G. (2015). Robustness and linear contracts. *American Economic Review, 105*, 536–563.

De Castro, L. I., & Yannelis, N. C. (2011). *Ambiguity aversion solves the conflict between efficiency and incentive compatibility*. University of Manchester Economic Discussion Papers Series, EDP-1106.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics, 75*, 643–669.

Gilboa, I., & Schmeidler, I. D. (1989). Maximin expected utility with non-unique prior. *Journal of Mathematical Economics, 18*, 141–153.

Grant, S. H., Kline, J. J., & Quiggin, J. (2006). *Lost in translation: honest misunderstandings and expost disputes*. Bond Business School Publications, Paper 3. http://epublications.bond.edu.au/business_pubs/3

Gul, F., & Postlewaite, A. (1992). Asymptotic efficiency in large exchange economies with asymmetric information. *Econometrica, 60*, 1273–1292.

Hahn, G., & Yannelis, N. C. (1997). Efficiency and incentive compatibility in differential information economies. *Economic Theory, 10*, 383–411.

Hart, O., & Moore, J. (1999). Foundations of incomplete contracts. *Review of Economic Studies, 66*, 115–139.

He, W., & Yannelis, N. C. (2015). Equilibrium theory under ambiguity. *Journal of Mathematical Economics, 61*, 86–95.

Holmström, B., & Myerson, R. B. (1983). Efficient and durable decision rules with incomplete information. *Econometrica, 51*, 1799–1819.

Knight, F. H. (1921). *Risk, uncertainty, and profits*. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Co.

Koutsougeras, L., & Yannelis, N. C. (1993). Incentive compatibility and information superiority of the core of an economy with differential information. *Economic Theory, 3*, 195–216.

Krasa, S., & Yannelis, N. C. (1994). The value allocation of an economy with differential information. *Econometrica, 62*, 881–900.

McLean, R., & Postlewaite, A. (2002). Informational size and incentive compatibility. *Econometrica, 70*, 2421–2454.

Mukerji, S. (1998). Ambiguity aversion and incompleteness of contractual form. *American Economic Review, 88*, 1207–1231.

Pulford, B. D. (2009). Is luck on my side: Optimism, pessimism, and ambiguity aversion. *Quarterly Journal of Experimental Psychology, 62*, 1079–1087.

Radner, R. (1968). Competitive equilibrium under uncertainty. *Econometrica, 36*, 31–58.

Radner, R. (1982). Equilibrium under uncertainty. In K. J. Arrow & M. D. Intriligator (Eds.), *Handbook of mathematical economics, vol. II*. Amsterdam: North Holland.

Savage, L. J. (1954). *Foundations of statistics*. New York, NY, USA: Wiley.

Segal, I. (1999). Complexity and renegotiation: A foundation for incomplete contracts. *Review of Economic Studies, 66*, 57–82.

Simon, H. A. (1956). *Reply: Surrogates for uncertain decision problems*. Office of Naval Research.

Wald, A. (1950). *Statistical decision functions*. New York, NY, USA: John Wiley & Sons.

Yannelis, N. C. (1991). The core of an economy with differential information. *Economic Theory, 1*, 183–197.

http://www.springer.com/978-981-10-1039-2

Social Norms, Bounded Rationality and Optimal
Contracts
Basov, S.
2016, IX, 176 p., Hardcover
ISBN: 978-981-10-1039-2