8 The Case of Incomplete Markets: Relating Risk Premiums to Economic Fundamentals

As already mentioned, the risk neutral valuation (RNV) method is too good to be true: It is not reasonable to consider that any risky asset can be priced without some knowledge of economic fundamentals such as supply and demand for consumption goods, investment goods and savings. But to incorporate such fundamentals one needs to model the behavior of consumers and investors toward risks. This chapter starts by presenting the expected utility criterion, which solves the famous St Petersburg paradox. This criterion can be very useful in several circumstances: modeling insurance decisions by individuals (Section 8.2) or modeling equilibrium prices in the markets where risks are exchanged (Section 8.3). However, the equilibrium approach to the assessment of risk premiums also has its limits, which are discussed in Section 8.4.

8.1 SOLVING THE ST. PETERSBURG PARADOX

This chapter examines the methods used by economists to value risks in isolation—that is, outside the magic world of complete markets where valuations can be deduced from market prices. The difficulties of valuing risks outside this magic world are well-illustrated by the St. Petersburg paradox, proposed in the eighteenth century by Nicolas Bernoulli and solved by his cousin Daniel Bernoulli. Consider the following lottery: Draw a coin several times, until the first head is obtained. If n draws have been necessary (which means that $[n-1]$ successive tails were obtained before the first head), then the lottery gives a prize of 2^n ducats. The question asked by Bernoulli is: How much should a rational decision maker value the right to participate in such a lottery? A little introspection shows that this number should not be much higher than a few ducats. However, because the probability of winning 2^n ducats is exactly $\frac{1}{2^n}$, the expected value[1] is infinite:

$$\mathcal{E}(\text{ gain }) = \left(\frac{1}{2}\times 2\right) + \left(\frac{1}{4}\times 4\right) + \left(\frac{1}{8}\times 8\right) + \cdots = +\infty.$$

Thus, there seems to be contradiction between the most natural criterion that can be used to value a lottery (the expectation or actuarial value of the gains) and what common sense would recommend.

Among the different solutions imagined to solve this paradox, the most reasonable was offered by Daniel Bernoulli, and axiomatized much later by Von Neumann and Morgenstern. It consists in computing the expected **utility** of gain, a method we now explain.

Bernoulli proposed to solve the St. Petersburg paradox by introducing psychological aspects. He claimed that the (marginal) satisfaction of winning one more ducat was inversely proportional to the amount already won, leading him to measure the total satisfaction (or utility as it is now called) of winning x ducats by a quantity that is proportional to the integral:

$$\int_1^x \frac{dt}{t},$$

which is equal to $\log x$, the logarithm of x.

With a logarithmic utility, the maximum price one should be ready to pay for participating in the St Petersburg lottery is exactly 4 ducats. Indeed, the expected utility of this lottery is equal to:[2]

$$\frac{1}{2}\log 2 + \frac{1}{4}\log 4 + \cdots \frac{1}{2^n}\log 2^n + \cdots = \log 4.$$

Now, if $u(x) = \log x$, we can write

$$\mathcal{E}u(\text{gain}) = u(4).$$

Thus, an investor characterized by a logarithmic utility has the same utility by participating in the St. Petersburg lottery or by holding exactly 4 ducats with probability 1.

In 1944, Von Neumann and Morgenstern developed an axiomatic justification of a more general criterion, which allows one to account for the specific characteristics of the decision maker. Von Neumann and Morgenstern proposed to allow any increasing function $u(x)$ to represent the utility of winning x ducats (or dollars to make it more modern). $u(.)$ is called the Von Neumann Morgenstern (in short, VNM) utility function of the decision maker.

It turns out that when u is concave, the expected utility of any lottery $\tilde{\ell}$ is always smaller than the utility of the expectation of this lottery:

$$u \quad \text{concave} \Rightarrow \mathcal{E}u(\tilde{\ell}) \leq u(\mathcal{E}\tilde{\ell}).$$

This is the risk aversion property: A decision maker with preferences associated with a concave VNM utility function always gains by replacing a lottery $\tilde{\ell}$ by

BOX 8.1 The Expected Utility Criterion

To any increasing function, u, one can associate a ranking $\succ$ among lotteries (represented by random variables $\tilde{\ell}_1, \tilde{\ell}_2$ characterizing the associated gains):

$$\tilde{\ell}_1 \succ \tilde{\ell}_2 \Leftrightarrow \mathcal{E}u(\tilde{\ell}_1) \geq \mathcal{E}u(\tilde{\ell}_2)$$

Lottery $\tilde{\ell}_1$ is preferred to lottery $\tilde{\ell}_2$ if and only if the expected utility of $\tilde{\ell}_1$ is greater than the expected utility of $\tilde{\ell}_2$.

> **BOX 8.2 The Risk Aversion Property**
>
> For any concave increasing VNM utility function, u, and any lottery, $\tilde{\ell}$, with a finite expectation $\mathcal{E}\tilde{\ell}$, the expected utility is smaller or equal to the utility of the expectation:
>
> $$\mathcal{E}u(\tilde{\ell}) \leq u(\mathcal{E}\tilde{\ell}).$$
>
> This means that the decision maker prefers a deterministic payment equal to $\mathcal{E}\tilde{\ell}$ to the stochastic payment $\tilde{\ell}$.

a non-random payment equal to the expectation of $\tilde{\ell}$. Without risk aversion of policyholders, insurance would not be a viable economic activity. Suppose, indeed, that the lottery $\tilde{\ell}$ represents the difference between an individual's wealth W and some insurable loss $\tilde{x}$:

$$\tilde{\ell} = W - \tilde{x}.$$

The expectation of $\tilde{\ell}$ equals the difference between the policyholder's wealth W and the actuarial premium $P = \mathcal{E}(\tilde{x})$ associated with the insurance of $\tilde{x}$. The insurance of this loss is a viable economic activity when the policyholder is ready to pay at least P to get rid of $\tilde{x}$:

$$u(W - \mathcal{E}(\tilde{x})) \geq \mathcal{E}u(W - \tilde{x}).$$

This is only guaranteed when u is concave.

In the polar case where u is convex, the reverse inequality is satisfied: for any convex function u and any lottery $\tilde{\ell}$, we have that

$$\mathcal{E}u(\tilde{\ell}) \geq u(\mathcal{E}\tilde{\ell}).$$

In this case, the decision maker is eager to take fair bets—that is, prefers any lottery $\tilde{\ell}$ to a non-random payment equal to the expectation of $\tilde{\ell}$.

8.2 CERTAINTY EQUIVALENT

For a given lottery (characterized by a random gain $\tilde{\ell}$) and a given decision maker (characterized by a VNM utility function $u(\cdot)$), the certainty equivalent (CE) is defined as the amount of money that gives the same utility to the decision maker (DM in the sequel) as participating in the lottery:

$$u(CE) = \mathcal{E}u(\tilde{\ell}).$$

In the example of the St. Petersburg lottery and the logarithm utility ($u(x) = \ln x$), the certainty equivalent is easy to obtain because we have seen in Section 8.1 (*see* footnote 2) that $\mathcal{E}\ln(\tilde{x}) = \ln 4$. Thus,

$$\ln(CE) = \ln 4,$$

which gives $CE = 4$. An agent having a logarithmic utility is ready to pay up to 4 ducats for participating in the St Petersburg lottery.

Other results are obtained with alternative specifications of u. One of the most popular of these alternative specifications is the **exponential utility:** $u(x) = t\left[1 - \exp -\frac{x}{t}\right]$, where $t > 0$ is a parameter that characterizes the risk tolerance of the decision maker. It is very popular because, when coupled with normality assumption on $\tilde{\ell}$, it gives raise to simple, explicit formulas. Indeed, when $\tilde{\ell}$ follows a normal distribution with mean μ and variance σ^2, one obtains:

$$\mathcal{E}\left(\exp -\frac{\tilde{\ell}}{t}\right) = \exp\left(-\frac{\mu}{t} + \frac{\sigma^2}{2t^2}\right).$$

Therefore:

$$\mathcal{E}u(\tilde{\ell}) = t\left[1 - \mathcal{E}\left(\exp -\frac{\tilde{\ell}}{t}\right)\right] = t\left[1 - \exp\left(-\frac{\mu}{t} + \frac{\sigma^2}{2t^2}\right)\right]$$
$$= u\left(\mu - \frac{\sigma^2}{2t}\right).$$

This gives a simple formula for the certainty equivalent:

$$CE = \left(\mu - \frac{\sigma^2}{2t}\right).$$

This formula is equivalent to the mean variance criterion presented in the Chapter 9[3]. However the expected utility criterion is more flexible, because it can cover situations where the risk tolerance index of the decision maker varies with his wealth or with other parameters that may change in relation with the risks incurred. For firms, this will lead to the notion of shareholder value function, which is presented in Chapter 12 of this book. The VNM utility function is the equivalent for individuals of the shareholder value function for corporations. We now present an application of the VNM utility function to insurance decisions.

APPLICATION: TO BUY OF NOT TO BUY INSURANCE?

Consider an individual or a corporation confronted with a risk of fire: With some probability p, a fire can occur, causing loss L. This decision maker can fully insure this risk by paying a premium P, but he can also retain the risk. In the first case, his utility is $u(W - P)$, where W denotes the policyholder's initial wealth. In the second case, his expected utility is

$$\mathcal{E}u = pu(W - L) + (1 - p)u(W).$$

In the absence of insurance, the certainty equivalent of the losses is defined implicitly by

$$u(W - CE) = pu(W - L) + (1 - p)u(W). \tag{8.1}$$

The rational decision is to buy insurance whenever $P < CE$, which can be interpreted as the maximum acceptable premium.

Numerical Example:

probability of fire	$p = 0.5\%$
loss in case of fire	$L = \$160,000$
individual's wealth	$W = \$250,000$
utility function	$u(x) = \sqrt{x}$.

We have by definition (8.1):

$$u(W - CE) = \sqrt{W - CE} = p\sqrt{W - L} + (1 - p)\sqrt{W}$$
$$= 0.005\sqrt{90,000} + 0.995\sqrt{250,000} = 499.$$

Thus,

$$CE = 250,000 - (499)^2 = 999.$$

The maximum premium that is acceptable for the customer is thus $CE = 999$.

Note that because of risk aversion (which is captured by the property that $u(\ell) = \sqrt{\ell}$ is concave), the actuarial premium $P = 0.005 \times (160,000) = 800$ is lower than CE. Therefore, any insurance premium P in the interval [800, 999] will be "viable" for the insurer, in the sense that it will attract the customer and simultaneously cover the expected losses of the insurer.

8.3 MARKETS FOR EXCHANGING RISKS

With the development of financial markets, the risks specialists of large corporations now have access to several market solutions for exchanging their risks, particularly the large ones. This section starts by a description of some of these market solutions and then relates their functioning to the theoretical model of risk exchange developed by Arrow and Borch.

8.3.1 The Practice

A first example of a market solution for exchanging risk is the Catastrophe Risk Exchange (CATEX) created in New York in 1996 (*see* Box 8.3). Others are derivative products like catastrophe options and catastrophe bonds traded in the Chicago Board of Trade (CBOT).

The CBOT has a long history in the derivatives industry. It was created in 1848 to provide an effective mechanism for buying and selling physical agricultural commodities. In 1992, just after Hurricane Andrew, the insurance and re-insurance industry was looking for new ways of managing and covering catastrophe risks. Several organizations, including the CBOT created new products such as cat options and cat bonds. These products allowed significant expansion of the capacity of the re-insurance market.

BOX 8.3 CATEX

Catastrophe Risk Exchange (CATEX) is a technology solution provider who has been authorized (and is regulated) by the New York Insurance Department as a re-insurance intermediary since 1996.

CATEX provides insurance and re-insurance firms with an instrument to manage more efficiently their portfolio of risks by offering an Internet-based transaction system to exchange catastrophe risks. CATEX is a sort of electronic marketplace on which insurance companies can list risks that they are eager to cede or to swap against other risks. In the first case, the format is a classical re-insurance treaty, in the second, it is a re-insurance swap transaction.

The standard unit of risk is \$ 1 million. As an example, an insurance company will increase the diversification on its portfolio by exchanging 15 units of Florida windstorm risk against 20 units of California earthquake risk.

Beside an enhanced management of catastrophe risks, CATEX also reduces transaction costs through its standardized technology solution that improves the administration of re-insurance treaties. This is a significant added value to the market, as reduced costs of administration may develop the number of transactions and create a more efficient market.

CATEX is not the only organization providing this type of services: the Bermuda Commodity Exchange has been developed with the same objective.

In 1995, CBOT proposed a new contract based on an index for insured loss developed by the Property Claims Service (PCS). These cat options are available for nine regions and states in the United States: National, East, Southeast, Northeast, Midwest, West Florida, Texas, and California.

BOX 8.4 Catastrophe Bonds

They allow the securitization of liabilities linked to catastrophe risks:

- A special purpose vehicle (SPV) provides re-insurance to the insurer (excess loss over some threshold ℓ^*).
- SPV issues a bond (called a catastrophe bond). The payment promised by this bond is indexed on the losses ℓ of the insurer.[4]
  - If the losses ℓ of the insurer are smaller than some predetermined threshold ℓ^*: SPV repays the full value of the bonds B.
  - If $\ell > \ell^*$: the SPV only pays $\ell - \ell^*$ to the insurer and $B - (\ell - L^*)$ to bondholders.

In this way, re-insurance risk is fully transferred to bondholders.

Insurance and re-insurance companies can buy the cash settled options to complement traditional re-insurance treaties; PCS options are traded on catastrophes trigger to protect insurers against catastrophe above a certain level.

8.3.2 The Theory

The actuary Karl Borch and the economist Kenneth Arrow have investigated the theoretical question of how risks should be shared within a community of individuals. This is an interesting application of the expected utility framework. Consider, for example, two individuals, indexed $i = 1,2$ and characterized by *VNM* utility functions $u_1(\cdot)$ and $u_2(\cdot)$. To simplify notation, suppose that there are only two states of the world $\{1,2\}$, with probabilities p_1, p_2.

Individuals' wealths are W_1 and W_2, but individual 1 loses L_1 in state 1 and individual 2 loses L_2 in state 2.

Suppose now that risks can be traded on an exchange at prices P_1, P_2. This means that individual i can receive θ_s \$ in state s in exchange for an unconditional payment $P_s\theta_s$ \$ (in both states).

For an insurance position θ_1, θ_2, individual i obtains the following expected utility:

$$U_i = p_1 u_i\left(W_{i1}^0 + \theta_1 - P_1\theta_1 - P_2\theta_2\right) + p_2 u_i\left(W_{i2}^0 + \theta_2 - P_1\theta_1 - P_2\theta_2\right),$$

where W_{i1}^0 and W_{i2}^0 are the "before trade" wealths given in Table 8.1. The optimal insurance position θ_1, θ_2 is obtained by maximizing U_i. It is characterized by the first-order conditions (zero derivatives):

$$\frac{\partial U_i}{\partial \theta_1} = p_1 u_i'(W_{i1})[1 - P_1] - p_2 u_i'(W_{i2})P_1 = 0,$$

and

$$\frac{\partial U_i}{\partial \theta_2} = -p_1 u_i'(W_{i1})P_2 + p_2 u_i'(W_{i2})[1 - P_2] = 0,$$

where W_{i1} and W_{i2} represent the "after-trade" wealths of individual i:

$$W_{is} = W_{is}^0 + \theta_s - P_1\theta_1 - P_2\theta_2.$$

Solving for P_1 and P_2, we obtain:

$$P_1 = \frac{p_1 u_i'(W_{i1})}{p_1 u_i'(W_{i1}) + p_2 u_i'(W_{i2})} \text{ for all } i,$$

TABLE 8.1. *Before Trade Wealths*

States \ Wealths	Individual 1	Individual 2
State 1	$W_1 - L_1$	W_2
State 2	W_1	$W_2 - L_2$

and similarly,

$$P_2 = \frac{p_2 u_i'(W_{i2})}{p_1 u_i'(W_{i1}) + p_2 u_i'(W_{i2})} \text{ for all } i,$$

Thus, the optimal position taken by each individual on the risk exchange is such that the marginal rate of substitution between incomes in the two states is equal to the ratio of prices:

$$\frac{p_1 u_i'(W_{i1})}{p_2 u_i'(W_{i2})} = \frac{P_1}{P_2} \quad \text{for} \quad i = 1, 2, \tag{8.2}$$

where W_is denotes the after-trade wealth of individual i in state s. The intuition for this can be established by contradiction: if $\frac{P_1}{P_2}$ was, for example, strictly less than this marginal rate of substitution, individual i could increase his expected utility by buying a little more of insurance against risk 1 and a little less of insurance against risk 2.

Equation (8.2) implies a very important property: after-trade wealths of risk-averse individuals are necessarily **co-monotonic**—that is, they move in the same way in the different states. If, for example, $W_{11} > W_{12}$ (individual 1's after trade wealth is higher in state 1) then by concavity of u_1 we have that $u_1'(W_{11}) < u_1'(W_{12})$ (marginal utility is higher in state 2), and therefore, by the above property:

$$\frac{u_1'(W_{11})}{u_1'(W_{12})} = \frac{p_2 P_1}{p_1 P_2} < 1.$$

Now the analogous inequality is also valid for individual 2:

$$\frac{u_2'(W_{21})}{u_2'(W_{22})} = \frac{p_2 P_1}{p_1 P_2} < 1.$$

This implies (by concavity of u_2) that $W_{21} > W_{22}$: individual 2's after-trade wealth is also higher in state 2.

8.3.3 Diversifiable Risk

Consider first the case where the possible losses of the two individuals are equal. This means $L_1 = L_2 = L$, so that there is no aggregate risk. Total wealth is constant across states:

$$W_{11} + W_{21} = W_{12} + W_{22} = W_1 + W_2 - L.$$

The co-monotonicity property established above implies that after-trade wealths of individuals must be constant across states: if, say, W_{11} was strictly less than W_{12}, this would be true also for individual 2 and by aggregation $W_{11} + W_{21} < W_{12} + W_{22}$. This would contradict the property that total wealth is constant across states.

After-trade incomes of individual 1 are thus

$$W_{11} = W_1 - L + \theta_1 - \theta_1 P_1 + \theta_2 P_2 \qquad \text{in state 1,}$$

TABLE 8.2. *After Trade Wealths When Risks Are Diversifiable*

States \ Wealths	Individual 1	Individual 2
State 1	$W_1 - p_1 L$	$W_2 - p_2 L$
State 2	$W_1 - p_1 L$	$W_2 - p_2 L$

and

$$W_{12} = W_1 - \theta_2 - \theta_1 P_1 + \theta_2 P_2 \qquad \text{in state 2.}$$

Equality of these two incomes implies that

$$\theta_1 + \theta_2 = L.$$

Moreover, buying one unit of each contract gives $1 for sure, which implies that

$$P_1 + P_2 = 1.$$

Now, using that $W_{11} = W_{12}$, equation (8.2) gives

$$\frac{P_1}{P_2} = \frac{p_1}{p_2}.$$

Thus, the ratio of prices $\frac{P_1}{P_2}$ is equal to the ratio of probabilities $\frac{p_1}{p_2}$. Because $P_1 + P_2 = p_1 + p_2 = 1$, it must be that $P_1 = p_1$ and $P_2 = p_2$. We see that in this case (the case of fully diversifiable risks) market prices of risks are just equal to their probabilities. Moreover, using all these properties, we can compute after-trade wealths:

$$W_{11} = W_{12} = W_1 - p_1 L,$$

and similarly,

$$W_{21} = W_{22} = W_2 - p_2 L.$$

This arrangement can be obtained by an insurance mutuality where individual risks are pooled. Each individual is completely insured, in exchange for the payment of an actuarial premium, $p_1 L$ for individual 1, and $p_2 L$ for individual 2 (note that the mutuality principle does not imply that premiums are necessarily equal across individuals: if $p_1 \neq p_2$, then premiums are different).

8.3.4 Aggregate Risks

In more general cases, perfect diversification cannot be attained: some aggregate risk remains. Assume, for example, that the loss L_1 of individual 1 in state 1 is smaller than the loss L_2 of individual 2 in state 2. Thus, aggregate wealth is smaller in state 2 (which we interpret as a recession) than in state 1 (which we interpret as a boom).

In this case, the co-monotonicity property implies that after-trade wealths of both individuals are also smaller in state 2 than in state 1:

$$W_{11} > W_{12} \qquad \text{and} \qquad W_{21} > W_{22}.$$

Then prices of risks are not equal to their probabilities:

$$\frac{P_1}{P_2} = \frac{p_1}{p_2}\frac{u_i'(W_{i1})}{u_i'(W_{i2})} < \frac{p_1}{p_2}.$$

Because $P_1 + P_2 = p_1 + p_2 = 1$, this implies that $P_1 < p_1$ and $P_2 > p_2$.

The market price of risk 1, which is also equal to the risk adjusted probability of state 1, is lower than the historical probability p_1 of state 1 (the reverse is true for state 2).

Property 8.1. *When there is some aggregate risk, the risk neutral probability measure*[5] (P_1, P_2) *gives more weight to unfavorable events (like a recession) and less weight to favorable events (like a boom) than the "historical" probability measure* (p_1, p_2). *In other words, the risk neutral measure is pessimistic.*

Finally, optimal risk sharing implies that every individual bears some fraction of aggregate risk: no one is perfectly insured. More precisely, the fraction of risk borne by each individual has to be proportional to the risk tolerance of the individual. To get simple formulas, we consider the case where utility functions are exponential:

$$u_i(x) = t_i\left[1 - \exp -\frac{x}{t_i}\right],$$

where $t_i > 0$ is the risk tolerance factor of individual i. Marginal utilities are given by $u_i'(x) = \exp -\frac{x}{t_i}$.

In this case, after-trade wealths W_{is} can be computed explicitly, because marginal rates of substitution are easy to compute

$$\frac{u_i'(W_{i1})}{u_i'(W_{i2})} = \exp\left[\frac{W_{i2} - W_{i1}}{t_i}\right].$$

These marginal rates of substitution must be equal across individuals. Thus, we must have:

$$\frac{W_{11} - W_{12}}{t_1} = \frac{W_{21} - W_{22}}{t_2}.$$

This means that the loss incurred by individual i in case of a recession (state 2) is proportional to t_i. Because the total loss $(W_{11} + W_{21}) - (W_{12} + W_{22})$ is equal to $L_2 - L_1$, we see that

$$W_{11} - W_{12} = \frac{t_1}{t_1 + t_2}(L_2 - L_1),$$

$$\text{and} \qquad W_{21} - W_{22} = \frac{t_2}{t_1 + t_2}(L_2 - L_1).$$

TABLE 8.3. *Aggregate Risks, After Trade Wealths*

States \ Incomes	Individual 1	Individual 2
State 1	$W_1 - \frac{t_1 L_1}{t_1+t_2}$	$W_2 - \frac{t_2 L_1}{t_1+t_2}$
State 2	$W_1 - \frac{t_1 L_2}{t_1+t_2}$	$W_2 - \frac{t_2 L_2}{t_1+t_2}$

Property 8.2. *Optimal risk sharing implies that aggregate risk is allocated to each individual in proportion to his risk tolerance.*

The efficient allocation of risk can be obtained through a co-insurance arrangement, whereby individual 1 transfers a fraction $\frac{t_2}{t_1+t_2}$ of his risk to individual 2 and accepts a fraction $\frac{t_1}{t_1+t_2}$ of individual 2's risk:

$$\begin{cases} W_{11} = W_1 - \frac{t_1 L_1}{t_1+t_2} \\ W_{12} = W_1 - \frac{t_1 L_2}{t_1+t_2} \end{cases}$$

Note that after-trade wealths of all individuals are lower in state 2 than in state 1 (this is because $L_2 > L_1$); aggregate risk is borne by all. However, if $t_1 < t_2$, individual 1 is less risk-tolerant than individual 2 and thus bears less risk.

Rule 8.1. *Optimal Risk Sharing*
The expected utility criterion gives a useful methodology for deciding how to use market instruments for sharing risks among firms or individuals. Two very different cases have to be distinguished: diversifiable risks and aggregate risks.

A risk is diversifiable if it does not impact the total resources of the community. This type of risk can be completely eliminated by mutualization: each individual fully transfers his own risk to the exchange, and prices of risks only reflect their actuarial values (no risk premium). For this type of risk, the risk neutral measure coincides with the historical probability measure.

By contrast, an aggregate risk (like the risk of a recession) has a non-predictible impact on the total resources of the community. This type of risk has to be shared by all the members of the community in proportion to their risk tolerance. Individuals or firms with an exposure that is larger than their risk tolerance (relative to the total risk tolerance of the community) are net sellers of the risk on the exchange, whereas individuals or firms with an exposure that is lower than their risk tolerance are net buyers of the risk.

8.4 THE LIMITS OF THE EQUILIBRIUM APPROACH

8.4.1 The Case of Incomplete Markets

When financial markets are complete, risks are optimally allocated, and RNV works. A sophisticated investor who knows the economic fundamentals can, in principle, compute the risk adjusted probability measure: multiplying historical probabilities by marginal utilities of consumption and scaling them to get a probability measure. But, of course, it is much simpler to deduce the risk-adjusted

measure from observed assets prices: this is the RNV methodology, which is equivalent to, but much simpler than, the equilibrium approach.

Everything breaks down when markets are not complete: risk allocation is not necessarily optimal and risk premiums cannot be deduced from the mere observation of asset prices. They depend in a complicated way on economic fundamentals. Chapter 9 explores the only case where risk premiums are easy to compute, even when markets are incomplete: if the joint distribution of all asset returns is Normal, the Capital Asset Pricing Model (CAPM) is indeed valid, and risk premiums are proportional to one simple factor: the regression coefficient of the asset return on the market return (the beta of the asset). Unfortunately, outside this Normal world, there is no simple way to compute risk premiums.

Of course, it could be argued that markets are "approximately" complete, in which case the RNV approach would be "approximately" valid. Unfortunately, there is a lot of indirect empirical evidence suggesting that financial markets are far from being complete. For example, market completeness would imply optimal risk sharing, and thus co-monotonicity of after-trade incomes. Recent crises tend to exhibit a different pattern, whereby some investors make a lot of money, whereas the majority of others loses: co-monotonicity is violated. More generally, many economists consider that financial markets sometimes magnify real shocks instead of dampening them, which is also a clear symptom of market incompleteness.

8.4.2 Knightian Uncertainty and the Ellsberg Paradox

The expected utility criterion has an important drawback: it does not fit very well the observed behavior of individuals who are uncertain about the (historical) probability distributions of risks.

The psychologist Ellsberg is famous for organizing the following experiment: individuals are asked to select between two urns containing black and white balls and then to draw one ball from the urn they have selected. The first urn contains the same numbers of black and white balls. The content of the second urn is uncertain. In the first experiment, individuals are told that they will receive a prize if they pick up a black ball. In the second experiment, it is the reverse: they receive a prize if they pick up a white ball. Asked to select one urn in each experiment, a great majority of individuals select the first urn (the one with a 50% chance to win) in both cases.

This seems natural but happens to be totally incompatible with the expected utility framework. Indeed, in this framework, each individual is supposed to form an expectation, p, of the frequency of black balls in the second urn. An individual who behaves like an expected utility maximizer will always select different urns in the two experiments (independently of his *VNM* utility function). If $p > 1/2$, he should select the first urn in the first experiment and the second urn in the second experiment. If $p < 1/2$, the converse is true. When $p = 1/2$, he should be strictly indifferent. None of these cases can explain the observed behavior of the majority of individuals. This observed behavior can only be explained by uncertainty aversion: most individuals prefer a situation where the distribution or risks is known to a situation where the distribution of risks is uncertain.

To capture the intuition behind Ellsberg's results, assume, for example, that individuals hesitate between three scenarios:

Scenario 1: The second urn also contains the same numbers of black balls and white balls.
Scenario 2: The second urn contains more black balls.
Scenario 3: The second urn contains more white balls.

Suppose now that the individual always considers the **worst-case scenario:** he aims at maximizing the **minimum** of the probabilities of gain in the three scenarios. It is clear then that he will select the first urn (which guarantees a probability of gain of 50%) in both experiments. This is because, in both experiments, there is a scenario that gives a strictly smaller probability of gain with the second urn.

8.4.3 When Markets Stop Functioning

The subprime crisis has also shown that some well-established markets can just stop functioning in some circumstances. This is the case in the example of the Asset-Backed Commercial Paper (ABCP) market that was a vital source of short-term financing for many firms and that completely dried up in the middle of the subprime crisis. One explanation often provided by economists is **adverse selection:** buyers had less information than sellers about the quality of the paper being sold. The following simple model, adapted from Akerlof's famous paper ("The Market for "Lemons": Quality Uncertainty and The Market Mechanism" *Quarterly Journal of Economics* 84(3), 488–500), illustrates how adverse selection can provoke the interruption of markets.

Consider a financial market where banks have the possibility to sell some of their assets to investors. The quality q of each asset, measured by the expected present value of its future cash flows, is known to the seller (the bank) but not to the buyers (the investors). Buyers only know the statistical distribution of the quality of asset being securitized. We are going to show that such markets are fragile: When adverse selection becomes too strong, these markets can stop functioning.

The maximum price P that buyers are ready to pay for each of these assets (that they cannot distinguish) is indeed the **average** quality of the assets on the markets:

$$P = \mathcal{E}(q|q \text{ is on the market}).$$

Suppose now that banks have access to a new investment opportunity, characterized by an expected present value of $(1+R)$ per unit of investment. If they can securitize one of their assets at price P and reinvest the proceeds into this new opportunity, they obtain an expected present value of $P(1+R)$, to be compared with q if they keep their original asset. Thus, banks are ready to sell their assets for a price below their quality (this is the economic justification of securitization), but there is a lower limit; the minimum price at which a bank accepts to securitize an asset of quality q is:

$$P_{\min}(q) = \frac{q}{1+R} < q.$$

If q was observable by buyers, then the market price would fully reflect it:

$$P = q,$$

and all gains from trade would be exploited:

$$P = q > P_{\min}(q) = \frac{q}{1+R} \text{ for all } q.$$

However, when q is not observable by buyers, the sellers might refuse to securitize some of their assets. This happens when

$$q > P(1+R).$$

By contrast, the banks are always willing to securitize their low-quality assets, such that:

$$q \leq P(1+R).$$

This is the essence of the adverse selection phenomenon: a (securitization) market always attracts the lowest quality sellers (banks). In general, trade is not efficient: the good quality assets are retained, and the banks that hold them cannot refinance. Sometimes the market can even stop altogether. To see this, assume that the statistical distribution of asset qualities is uniform on some interval $[q_0, q_1]$. When $q_1 - q_0$ is small (mild adverse selection), all assets are securitized (efficient trading) at price $P = \frac{1}{2}(q_0 + q_1)$. This happens if and only if $P \geq \frac{q_1}{1+R}$, which is equivalent to $q_1 \leq q_0 \frac{1+R}{1-R}$.

In the opposite case:

$$q_1 > q_0 \frac{1+R}{1-R},$$

only the assets with quality in the interval $[q_0, P(1+R)]$ are put on the market, and the securitization price P is equal to the average quality of the assets on the market:

$$P = \frac{1}{2}[q_0 + P(1+R)],$$

which allows to determine the equilibrium price:

$$P = \frac{q_0}{1-R},$$

and the volume of trade

$$V = P(1+R) - q_0 = \frac{2q_0 R}{1-R}.$$

Suppose now that the minimum quality of assets suddenly deteriorates: q_0 becomes zero. Then the volume of trade also becomes equal to zero: the market stops functioning altogether. The presence of very low-quality assets that cannot be distinguished from the others is enough to precipitate the complete halt of securitization markets.