RISK, BEHAVIOR AND EVOLUTION

Dissertation

zur Erlangung des Grades
Doktor der Wirtschaftswissenschaften (Dr. rer. pol.)
am Fachbereich Wirtschaftswissenschaften
der Universität Konstanz

vorgelegt von:
Nick Netzer
Mangoldstraße 11
78462 Konstanz

Konstanz, 24. Oktober 2007.

Tag der mündlichen Prüfung: 5. Februar 2008.
Prüfungskommission:
Prof. Dr. Urs Fischbacher (Vorsitzender), Universität Konstanz
Prof. Dr. Friedrich Breyer (1. Referent), Universität Konstanz
Prof. Dr. Carlos Alós-Ferrer (2. Referent), Universität Konstanz

Für Maria

Table of Contents

Introduction and Summary

This dissertation consists of four self-contained research papers that have been written during the time from October 2004 to October 2007, while I was participant of the "Doctoral Programme in Quantitative Economics and Finance" at the University of Konstanz. I have two major research interests. The first covers questions of optimal economic policy, with a special focus on optimal taxation and social insurance. The contributions in Chapters 1 and 2 belong to this area. My second field of interest concerns more basic questions of individual economic behavior and its foundations. Chapters 3 and 4 address such issues. This section provides a brief introduction to the following chapters and summarizes the main results.

Chapter 1 is the result of joint work with Florian Scheuer (Massachusetts Institute of Technology). Parts of it have been published under the title "Taxation, Insurance and Precautionary Labor" in the *Journal of Public Economics*, Vol. 91, 2007, p. 1519-1531. We examine optimal taxation and social insurance with adverse selection in competitive insurance markets. In a previous literature, it has been shown that, with perfect insurance markets, social insurance improves welfare since it is able to redistribute without creating distortions. This result has been taken as robust to the introduction of adverse selection as this would only provide additional justifications for social insurance. We show, however, that adverse selection can weaken the case for social insurance compared to a situation with perfect markets. Whenever social insurance mitigates private underinsurance, it also causes welfare-reducing effects by decreasing precautionary labor supply and hence tax revenue. In addition, adverse selection may substantially reduce the redistributive potential of social insurance. We illustrate our general results using different equilibrium concepts for the insurance market.

Chapter 2 addresses the question of optimal taxation if individual property rights are costly to enforce, due to potential crime. The level of crime is often claimed to be correlated with the degree of inequality in a society, and redistribution is suggested as a remedy. A producer-predator model of theft in large heterogeneous societies is analyzed, and the effect of public redistribution on individual incentives and consumption levels is examined. It turns out that redistribution of income by moving from a linear to a progressive tax schedule can be Pareto improving under a large variety of circumstances. If the ratio of high-skilled to low-skilled individuals is small, progressive taxation acts as a substitute for direct law enforcement. Otherwise, redistribution and law enforcement are complements, in the sense that the introduction of progressive taxation should go along with stricter law enforcement. Altogether, the existence of crime is never Pareto efficient as soon as progressive taxation is admitted.

Chapter 3 explores a general model of the evolution and adaption of hedonic utility. It is shown that optimal utility will be increasing strongly in regions where choices have to be made often and decision mistakes have a severe impact on fitness. Several applications are suggested. In the context of intertemporal preferences, the model offers an evolutionary explanation for strong short-run impatience. It also explains the existence of conflicting short- and long-run interests that lead to dynamic inconsistency. Concerning attitudes towards risk, an evolutionary explanation is given for S-shaped value functions that adjust to the decision-maker's environment.

Chapter 4 is the result of joint work with Carlos Alós-Ferrer (University of Konstanz). We develop a characterization of stochastically stable states for the logit-response learning dynamics in games, with arbitrary specification of revision opportunities. The result allows us to show convergence to the set of Nash equilibria in the class of best-response potential games and the failure of the dynamics to select potential maximizers beyond the class of exact potential games. We also study to which extent equilibrium selection is robust to the specification of revision opportunities. Our techniques can be extended and applied to a wide class of learning dynamics in games.

Einleitung und Zusammenfassung

Diese Doktorarbeit besteht aus vier eigenständigen Forschungsarbeiten, die in der Zeit von Oktober 2004 bis Oktober 2007 geschrieben wurden, während ich Teilnehmer am "Doctoral Programme in Quantitative Economics and Finance" an der Universität Konstanz war. Meine Forschungsinteressen erstrecken sich auf zwei Gebiete. Das erste beinhaltet Fragen nach der optimalen Wirtschaftspolitik, und hier insbesondere nach optimaler Besteuerung und Sozialversicherung. Die Beiträge in den Kapiteln 1 und 2 gehören in diesen Bereich. Zum anderen interessiere ich mich für grundlegende Fragen über individuelles ökonomisches Verhalten und dessen Fundierung. Die Kapitel 3 und 4 beschäftigen sich mit solchen Themen. Dieser Abschnitt stellt eine kurze Einführung in die folgenden Kapitel dar, und fasst die wichtigsten Ergebnisse zusammen.

Kapitel 1 ist das Ergebnis einer gemeinsamen Arbeit mit Florian Scheuer (Massachusetts Institute of Technology). Teile davon wurden unter dem Titel "Taxation, Insurance and Precautionary Labor" im *Journal of Public Economics*, Vol. 91, 2007, p. 1519-1531, veröffentlicht. Wir untersuchen optimale Besteuerung und Sozialversicherung wenn auf privaten Versicherungsmärkten adverse Selektion vorliegt. In der bisherigen Literatur wurde gezeigt, dass Sozialversicherung bei perfekten privaten Versicherungsmärkten wohlfahrtssteigernd wirkt, da sie Einkommen wie gewünscht umverteilt, ohne Verzerrungen auszulösen. Es wurde dann davon ausgegangen, dass dieses Ergebnis auch für Märkte unter asymmetrischer Information gelten müsse, da adverse Selektion bekanntermaßen selbst einen Rechtfertigungsgrund für Sozialversicherung darstellt. Wir zeigen jedoch, dass die Annahme von adverser Selektion das Argument für Sozialversicherung sogar abschwächen kann, im Vergleich zu perfekten Märkten. Wenn eine Sozialversicherung tatsächlich das Problem der privaten Unterversicherung behebt, dann reduziert sie gleichzeitig "Vorsichtsmotive" und somit das Arbeitsangebot, die Steuereinnahmen und die gesellschaftliche Wohlfahrt. Zudem kann adverse Selektion das Umverteilungspotential von Sozialversicherung erheblich reduzieren. Wir veranschaulichen unsere allgemeinen Ergebnisse anhand von verschiedenen Gleichgewichtskonzepten für Versicherungsmärkte.

Kapitel 2 beschäftigt sich mit optimaler Besteuerung unter Berücksichtigung der Tatsache, dass Eigentumsrechte aufgrund von Kriminalität nicht kostenlos durchgesetzt werden können. Das Niveau der Kriminalität wird in der öffentlichen Diskussion häufig mit dem Grad der materiellen Ungleichheit in einer Gesellschaft assoziiert, und Umverteilung wird als mögliche Lösung vorgeschlagen. Die vorliegende Arbeit untersucht die Bestimmungsgrößen von Kriminalität anhand eines "Producer-Predator" Modells einer großen, heterogenen Gesellschaft, und analysiert den Einfluss von Umverteilung auf individuelle Anreize und materiellen Wohlstand. Es stellt sich heraus, dass unter vielerlei Umständen der Übergang von einem linearen zu einem progressiven Steuertarif eine Pareto-Verbesserung herbeiführen kann. Wenn das Verhältnis von hoch- zu niedrigqualifizierten Personen klein ist, dann stellt die progressive Besteuerung ein Substitut für direkte Rechtsdurchsetzung dar. Anderenfalls sind progressive Besteuerung und Rechtsdurchsetzung Komplemente, da mit der Einführung der Progression die zur Rechtsdurchsetzung eingesetzten Steuermittel ausgeweitet werden sollten. Insgesamt ist es niemals effizient, die Existenz von Kriminalität zu akzeptieren sofern das Instrument der progressiven Besteuerung verfügbar ist.

Kapitel 3 untersucht ein allgemeines Modell der Evolution und Adaption von kardinalen Nutzenfunktionen. Es wird gezeigt, dass das evolutorisch optimale Nutzenniveau besonders stark in Konsumregionen ansteigt, in denen häufig Entscheidungen getroffen werden müssen und in denen Fehlentscheidungen einen besonders gravierenden Einfluss auf die biologische Fitness haben. Dieses Ergebnis wird auf verschiedene Anwendungen übertragen. Im Kontext von intertemporalen Präferenzen bietet das Modell eine evolutorische Begründung für eine ausgeprägte Ungeduld bezüglich kurzfristiger Wartezeiten. Es erklärt zudem die Existenz von eventuell miteinander in Konflikt stehenden kurz- und langfristigen Interessen, die zu dynamischer Inkonsistenz führen können. In Bezug auf Risikoeinstellungen ergibt sich eine Erklärung für S-förmige Wertfunktionen, die sich an die Umgebung des Entscheidungsträgers anpassen.

Kapitel 4 ist das Ergebnis einer gemeinsamen Arbeit mit Carlos Alós-Ferrer (Universität Konstanz). Wir entwickeln eine Charakterisierung der stochastisch stabilen Zustände der "Logit-Response" Lerndynamik für allgemeine Spiele. Der Prozess, der die Revisionsmöglichkeiten der Spieler bestimmt, muss dabei nicht den üblichen Beschränkungen unterliegen. Mit Hilfe unseres Ergebnisses ist es möglich zu zeigen, dass die Logit-Response Dynamik in der Klasse der "Best-Response" Potentialspiele zur Menge der Nash-Gleichgewichte konvergiert. Das Ergebnis, nach dem in (exakten) Potentialspielen diejenigen Strategienprofile stochastisch stabil sind, die die Potentialfunktion maximieren, gilt bereits in der Klasse der gewichteten Potentialspiele nicht mehr. Wir untersuchen zudem, inwiefern die Gleichgewichtsauswahl von der genauen Spezifikation der Revisionsmöglichkeiten abhängt. Die von uns entwickelte Methodik lässt sich erweitern und kann auf eine große Anzahl von verschiedenen Lerndynamiken angewandt werden.

Chapter 1

Taxation, Insurance and Precautionary Labor

1.1 Introduction and Literature

It is known since Mirrlees (1971) that the problem of taxation is fundamentally linked to asymmetric information between the government and workers. Only with the assumption that the government cannot observe individual productivities does the need for distortionary income taxation arise. The contributions of Rothschild and Stiglitz (1976) and Wilson (1977) have shown that a similar issue makes the problem of equilibrium in competitive insurance markets a relevant question. If insurance companies cannot observe individual risk types, the resulting market allocation will not in general be efficient. With this article, we aim at providing a theory that ties together these two branches of information economics and at highlighting previously ignored interactions between distortionary taxation, social insurance and imperfect insurance markets.

Our starting point are the existing models of taxation and social insurance, such as Rochet (1991), Cremer and Pestieau (1996) and Henriet and Rochet (2004). Using the assumption of perfect insurance markets, this literature has concluded that social insurance can be a useful instrument for redistribution as it evens out differences in private insurance premiums without causing distortions. This result has been taken as robust to the introduction of adverse selection, which would only constitute an additional justification for social insurance. Indeed, Wilson (1977) and Eckstein et al. (1985) have shown that the government might be able to Pareto improve upon the market allocation by introducing social insurance in an insurance market with adverse selection. Social insurance may mitigate the underinsurance of low risk types and at the same time reduce the average premium for high risk types, thus making all individuals better off. The simple intuition that equity and efficiency effects complement one another as motivations for social insurance has therefore prevailed.

In this chapter, we demonstrate that this reasoning is invalid. It ignores the interdependencies that emerge when the models of taxation and of insurance markets are combined thoroughly. The link between the two strands is a theory of precautionary labor supply that we develop in section 1.2.[1] While the theory of

[1]Our results are based on the insights of Kimball (1990). A detailed discussion of the literature on precautionary motives is relegated to section 1.2.

optimal income taxation requires labor supply to be endogenous, models of competitive insurance markets with adverse selection imply that not all uncertainty can be resolved. It is therefore necessary to derive the determinants of labor supply under uncertainty. We show that, under a broad range of reasonable assumptions, greater uncertainty leads individuals to increase their labor supply. In section 1.3, we use this result to demonstrate that adverse selection can weaken the case for social insurance compared to a situation with perfect markets. Social insurance might indeed work against the inefficiency of underinsurance. At the same time, however, individuals faced with less uncertainty will reduce labor supply, and tax revenues will decline. This negatively affects social welfare. Furthermore, with adverse selection it is no longer clear whether social insurance can redistribute income at all, as private premiums do not necessarily correspond to individual risks any more.

In section 1.4, we illustrate our general results using different equilibrium concepts for the insurance markets. Several insights can be drawn from these illustrations. Social insurance will alleviate the inefficiency of underinsurance in the Rothschild-Stiglitz framework but has negative effects in the labor market by reducing precautionary labor supply. This endorses the case for only partial social insurance or even for complete renunciation. If the equilibrium is of the Wilson pooling type, social insurance additionally loses its main potential for redistribution. In case of a second-best Miyazaki-Wilson equilibrium, no positive efficiency effects of social insurance remain, while it still entails labor supply distortions and suffers from reduced redistributive power. In sum, we conclude that the case for social insurance is weakened by the presence of adverse selection in private insurance markets irrespective of the specific equilibrium concept considered.

The most similar existing work is the contribution by Boadway et al. (2006) who were the first to examine optimal taxation with adverse selection and ex-post moral hazard in private insurance markets.[2] Based on Rothschild-Stiglitz separating equilibria they find the case for social insurance strengthened by market inefficiencies. However, while labor supply is chosen under uncertainty in our model, their results are based on the assumption that labor supply decisions take place *after* a

[2]We refrain from analyzing moral hazard in this article.

possible damage has been realized. This reduces the impact of underinsurance on individual decisions to income effects. We do not want to eliminate the precautionary effects resulting from labor supply under uncertainty which play a crucial role in understanding the interaction between taxation and insurance. In addition, as we will show, results based on the specific Rothschild-Stiglitz equilibrium concept are not a complete representation for the general case.

1.2 Labor Supply under Uncertainty

In this section, we derive important results on labor supply under uncertainty that will be used in our model of optimal taxation and social insurance. Making use of the insights of Kimball (1990), we establish a theory of "precautionary labor". While the theory of precautionary savings has received some attention (Sandmo (1970), Abel (1988), Kimball (1990)), the problem of labor supply under uncertainty is less explored. Eaton and Rosen (1979) and (1980), Hartwick (2000), Parker et al. (2005) and Floden (2006) consider the case of endogenous labor with wage uncertainty. Low and Maldoom (2004) show that progressive taxation can negatively affect labor supply because it reduces such wage risk and hence precautionary motives. We do not model wage risk but an income independent risk to consumption. Labor supply is chosen *before* the risk is realized. This gives rise to the question whether risk induces people to work more or less than they would in case of certainty.[3]

We restrict our attention to Bernoulli random variables that result from a possible damage D which occurs with probability p. This restriction allows us to generalize the results of Kimball (1990) to situations with preexisting risk, i.e. to examine the question how changes in uncertainty affect labor supply in situations that are already risky.[4] Let $\theta(\beta)$ denote such a Bernoulli random variable, where the parameter $\beta \in [0, 1]$ stands for the share of the damage that is insured. It can be used to vary both expected value $\text{E}[\theta(\beta)] = p\,(1-\beta)D$ and variance $\text{Var}[\theta(\beta)] = p\,(1-p)[(1-\beta)D]^2$ of the risk. Furthermore, as throughout the chapter, an additively separable utility function $U(c, L) = u(c) + v(L)$ is assumed, where c

[3]Similar effects appear in Stiglitz (1982). However, his contribution lacks the theoretical tools developed by Kimball (1990), which will allow for a very clear analysis of labor supply under uncertainty in the following.

[4]For a detailed discussion of "preexisting risk" and its implication for the applicability of standard measures of risk aversion see Ross (1981).

denotes consumption and L denotes labor supply.[5] The standard conditions $u' > 0$, $u'' < 0$, $v' < 0$ and $v'' < 0$ are assumed to hold. Denote the productivity of an individual by w. Firms can observe w and pay wages according to marginal productivity such that earned income is wL. The individual receives an additional, exogenous and state independent income T.

The first order condition for labor supply L^* that maximizes expected utility in the presence of a given consumption risk $\theta(\beta)$ is

$$w \, \mathrm{E}\left[u'\left(wL^* + T - \theta(\beta)\right)\right] = -v'(L^*), \tag{1.1}$$

where E is the expectations operator.[6] (1.1) is a standard condition stating that labor supply is determined so as to equalize expected marginal utility and disutility from work. To answer the question how risk affects labor supply, we examine the move from $\theta(\beta)$ to the risk $\theta(\beta_1) + (\beta_1 - \beta)pD$, which constitutes a change in variance, leaving the expected value unaffected. We define the corresponding *equivalent precautionary premium* $\Psi(\beta, \beta_1)$ for such a move implicitly as follows:[7]

$$\mathrm{E}\left[u'\left(wL^* + T - \theta(\beta) - \Psi(\beta, \beta_1)\right)\right] = \mathrm{E}\left[u'\left(wL^* + T - \theta(\beta_1) - (\beta_1 - \beta)pD\right)\right]. \tag{1.2}$$

It has the following interpretation: The compensated change in insurance will have the same effect on the LHS of (1.1) and therefore on optimal labor supply as a lump-sum reduction of income by $\Psi(\beta, \beta_1)$. Both affect the optimality condition in the same way. Therefore, statements about the effect of risk on labor supply can be restated as income effects triggered by a decrease of income by Ψ.

Implicit differentiation of (1.2) yields an explicit formulation for $\partial\Psi(\beta, \beta_1)/\partial\beta_1$. Most relevant is the evaluation of this derivative at $\beta_1 = \beta$, which gives the income change that would have the same effect on labor supply as a small change in

[5] We need the assumption of separability only to keep the exposition of our labor supply theory concise. As shown by Kimball (1990), the results can be transferred to the case of nonseparable utility. In the later sections, separability is not necessary either but only used since we resort to the theory of labor supply under uncertainty there. We also assume the function v to be at least twice and u at least three times continuously differentiable. Finally, we do not discuss possible corner solutions for labor supply or consumption.

[6] The sufficient second order condition for a maximum is satisfied.

[7] As shown by Kimball (1990), the discussed premium is simply the equivalent risk premium developed by Pratt (1964), applied to the first derivative of u.

insurance, starting from a situation with insurance β. We obtain after a few rearrangements

$$\left.\frac{\partial\Psi(\beta,\beta_1)}{\partial\beta_1}\right|_{\beta_1=\beta} = \left(-\frac{\Delta u''(.)/(1-\beta)D}{\mathrm{E}[u''(.)]}\right)\left(\frac{1}{2}\frac{\partial \mathrm{Var}[\theta(\beta)]}{\partial\beta}\right), \tag{1.3}$$

where $\Delta u''(.)$ is the difference of $u''(.)$ between consumption levels in case of no damage and damage, and $\mathrm{E}[u''(.)]$ is the expected value of $u''(.)$.[8]

The first term in brackets on the RHS of (1.3) is the *generalized coefficient of absolute prudence* $\eta^G(\beta)$. As β converges to 1, i.e. the examined situation converges to a situation without risk, the coefficient η^G converges to the prudence η as defined by Kimball (1990), which is simply the coefficient of absolute risk aversion for the function $u'(.)$, i.e. $\eta = -u'''/u''$. Opposed to the results by Kimball, our secant formulation of $\eta^G(\beta)$ makes it possible to examine situations with preexisting risk ($\beta < 1$). From (1.3) follow first implications for labor supply under uncertainty, which are summarized in the following lemma.

Lemma 1.1. *A marginal increase in insurance coverage which is compensated by an actuarially fair premium adjustment, decreases labor supply if and only if* $\eta^G(\beta) > 0$.

Proof. By (1.3) and $\partial\mathrm{Var}[\theta(\beta)]/\partial\beta < 0$, the compensated increase in insurance has the same effect as an increase in income if and only if $\eta^G(\beta) > 0$. With separable preferences, leisure is a normal good. Hence, an increase in income decreases labor supply. □

Conversely, a higher labor supply will be the reaction to less insurance iff $\eta^G > 0$; the individual has a motive for *precautionary labor*. The size of η^G indicates how strong this motive is. A sufficient condition for $\eta^G(\beta)$ to be positive is that $u'''(.)$ is positive in the relevant range of consumption levels. This in turn is a necessary condition for constant or decreasing risk aversion, both in absolute and relative terms.[9] Therefore, under the common and realistic assumption of non-increasing risk aversion, precautionary labor effects do exist. If not indicated otherwise, we will assume this for the rest of the chapter.

The results so far were derived for changes in risk that leave the expected damage unaffected. If this is not the case, additional income effects arise. It is still

[8] The derivation makes use of the fact that $\Psi(\beta,\beta) = 0$ holds.
[9] See Appendix 1.A.1 for a proof.

useful to distinguish between pure risk effects via the variance and income effects via expected values. Lemma 1.2 demonstrates this decomposition.

Lemma 1.2. *The total effect of a marginal increase in β on labor supply is*

$$\frac{\partial L^*}{\partial \beta} = \frac{\partial L^*}{\partial T}\left[pD - \frac{\partial \Psi}{\partial \beta}\right], \tag{1.4}$$

where $\partial L^*/\partial T < 0$ is the negative income effect and $\partial \Psi/\partial \beta$ stands short for the expression (1.3).[10]

Proof. See Appendix 1.A.2. □

First, higher coverage increases expected income by pD. This effect would vanish if an insurance premium were adjusted actuarially fairly. Second, the change in the variance has the same effect as a decrease of income by the premium Ψ that is raised by an increased insurance coverage β. As shown above, this premium will in general be negative.

1.3 Optimal Taxation and Social Insurance

1.3.1 The Model

This section derives conditions for optimal government policy in the presence of adverse selection in insurance markets. This is done without an explicit model of such market imperfections. We demonstrate in section 1.4 that different equilibrium concepts can easily be incorporated.

Our model setup is similar to Cremer and Pestieau (1996) and Boadway et al. (2006). We consider a society that consists of N individuals described by two characteristics: their productivity and their probability of incurring a damage of size D. There are W different productivity levels w_i, $i = 1, ..., W$ and two damage probabilities p_j, $j = L, H$, with $p_L < p_H$. In what follows, the index i will always refer to productivity while j refers to damage probability. We denote the proportion of individuals in the population that have productivity w_i and damage probability p_j by n_{ij}. The population average of the risk probability is then given by $\bar{p} = \sum_{i,j} n_{ij} p_j$. The average probability within productivity group i is $\bar{p}_i = (1/(n_{iL} + n_{iH})) \sum_j n_{ij} p_j$.

[10]Throughout the rest of the chapter, we stick to this convention, i.e. we write $\partial \Psi/\partial \beta$ short for the expression (1.3).

As commonly assumed in the theory of optimal taxation, the government maximizes the unweighted utilitarian objective.[11] It can neither observe individual productivities nor damage probabilities but only knows the joint distribution of both characteristics. Hours worked are unobservable as well, so that taxes have to be conditioned on observable labor income and will cause distortions. Unobservability of individual risks is not only a realistic assumption but introduces the case for social insurance in the first place. Otherwise, income taxation could be directly conditioned on risk as shown by Boadway et al. (2006).

The tax schedule is restricted to a constant marginal tax rate τ and a lump-sum transfer T.[12] In addition, the government can force the citizens to insure a share α of the possible damage. Such social insurance is financed by a uniform contribution $\bar{p}\,\alpha D$ by each individual. The remaining risk can be insured privately. The contract that individual ij purchases is denoted by $\mathcal{I}_{ij} = (\beta_{ij}, d_{ij})$, where β_{ij} is the privately insured share of the damage and d_{ij} is the premium.

The time structure is as follows. First, the government sets its policy $\mathcal{P} = (\tau, T, \alpha)$. Taking $\mathcal{P}$ as given, individuals simultaneously choose their labor supply and purchase their insurance contract. They also pay taxes and social insurance contributions and receive the transfer. Finally, the damage occurs according to the given probabilities. After possible payments of social and private insurance, consumption takes place.[13]

1.3.2 Optimal Government Policy

The timing of our model is such that individuals simultaneously choose their labor supply and a private insurance contract. For the sake of exposition, however, suppose for the moment that an individual's choice of insurance (β, d) is fixed exogenously. Then optimal labor supply $L^*_{ij}(\tau, T, \alpha, \beta, d)$ can be determined. It is implicitly defined by a standard first order condition that can be differentiated to

[11] Introducing individual weights in the welfare function would not fundamentally change the following results. It would simply induce additional weights on the "social valuations" to be defined below. Doing so would allow us to construct the whole second-best Pareto frontier.

[12] A straightforward extension would allow for non-linear taxation. However, we doubt that this changes the results significantly since the underlying intuition, as will be discussed below, is independent of the assumed linearity.

[13] The assumption that individuals have to decide on labor supply before the risk is realized is crucial. It is highly realistic for a large class of risks such as longevity, unemployment or illness. These risks are commonly covered by social insurance.

obtain comparative static effects as shown in the proof of Lemma 1.2. In our notation, the derivative of L^*_{ij} with respect to α already includes the effect of the increase in the social insurance contribution $\overline{p}\alpha D$. By contrast, the derivative with respect to β does not take into account a change in the premium. Where needed, the effect that accounts for such a change is marked with the letter A:[14]

$$\left.\frac{\partial L^*_{ij}}{\partial \beta}\right|_A = \frac{\partial L^*_{ij}}{\partial T}\left[p_j D - \frac{\partial d}{\partial \beta} - \frac{\partial \Psi_{ij}}{\partial \beta}\right]. \tag{1.5}$$

Substitution of L^*_{ij} into the expected utility function yields the indirect expected utility function $V^*_{ij}(\tau, T, \alpha, \beta, d)$.

The actual private insurance contracts are endogenous and will depend on the specific equilibrium concept as illustrated in Section 1.4. At this point it is only necessary to emphasize that they will depend on the policy $\mathcal{P}$, i.e. $\beta_{ij} = \beta_{ij}(\tau, T, \alpha)$ and $d_{ij} = d_{ij}(\tau, T, \alpha)$. For the purpose of comparative statics with respect to the policy parameters, it will be convenient, however, to express the premium d_{ij} as a differentiable function of the coverage β_{ij}, i.e. $d_{ij} = d_{ij}(\beta_{ij})$. This is indeed possible for all equilibrium concepts that we will consider later.[15] Functions that account for such type-specific equilibrium effects are marked by two asterisks, i.e. $L^{**}_{ij}(\tau, T, \alpha) = L^*_{ij}(\tau, T, \alpha, \beta_{ij}(\tau, T, \alpha), d_{ij}(\beta_{ij}(\tau, T, \alpha)))$. Indirect utility $V^{**}_{ij}(\tau, T, \alpha)$ is defined analogously.[16] With this notation, and assuming no exogenous revenue requirement, the government's optimization problem is

$$\max_{T,\tau,\alpha} \sum_{i,j} n_{ij}\, V^{**}_{ij}(\tau, T, \alpha) \quad \text{s.t.} \quad \sum_{i,j} n_{ij}\left(\tau w_i L^{**}_{ij}(\tau, T, \alpha) - T\right) = 0. \tag{1.6}$$

As common in the theory of optimal taxation, general explicit solutions to the optimization problem cannot be obtained, so that we resort to a detailed discussion of the necessary optimality conditions. Assuming interior solutions, we derive the first order conditions for problem (1.6) and transform them to obtain the following

[14] It is assumed in this formulation that the premium d can be derived from the coverage β through a differentiable function, so that $\partial d/\partial \beta$ is well-defined. This will be further explained below.

[15] The approach will be modified slightly in Section 1.4.3, to allow for the case that the premium does also depend on the coverage that other types obtain in equilibrium.

[16] Although each individual chooses its most preferred contract $\mathcal{I}_{ij}$ out of the set of available contracts, the function V^*_{ij} is not necessarily an optimal value function with respect to (β_{ij}, d_{ij}) in the sense of the Envelope theorem. This would be the case if, for instance, insurance markets were perfect and each individual purchased full coverage at an actuarially fair premium. With adverse selection, individuals will not generally be able to purchase such optimal contracts.

Proposition 1. We use three important concepts. The first is the "net social marginal valuation of an individual's income", b_{ij}, well-known from the theory of optimal taxation,

$$b_{ij} = \frac{1}{\gamma}\frac{\partial V_{ij}^*}{\partial T} + \tau w_i \frac{\partial L_{ij}^*}{\partial T}, \tag{1.7}$$

where γ is the Lagrange multiplier associated with the revenue constraint, whose optimal value equals the welfare value of a marginal increase in government revenues.[17] It captures the effect of an increased transfer T on the objective via the individual's utility and via the effect on the budget constraint through labor supply changes, both measured in terms of government revenues. On our model, government policy has additional effects on the objective via the insurance market equilibrium. Therefore, the concept of "net social marginal valuation of an individual's insurance", g_{ij}, is useful:

$$g_{ij} = \frac{1}{\gamma} \left.\frac{\partial V_{ij}^*}{\partial \beta}\right|_A + \tau w_i \left.\frac{\partial L_{ij}^*}{\partial \beta}\right|_A. \tag{1.8}$$

It captures the effect of a changing equilibrium contract via utility and via the budget constraint. Finally, the Slutzky decomposition

$$\frac{\partial L_{ij}^*}{\partial \tau} = -w_i L_{ij}^* \frac{\partial L_{ij}^*}{\partial T} - w_i \frac{\partial L_{ij}^c}{\partial w_i^n} = -w_i L_{ij}^* \frac{\partial L_{ij}^*}{\partial T} - \epsilon_{ij} \frac{L_{ij}^*}{1-\tau} \tag{1.9}$$

is used, where $w_i^n = (1-\tau)w_i$ denotes the net wage, L_{ij}^c is the Hicksian (compensated) labor supply function, and $\epsilon_{ij} = w_i^n / L_{ij}^* \times \partial L_{ij}^c / \partial w_i^n$ is the positive elasticity of L_{ij}^c with respect to the net wage. With this notation, we have

[17] If not stated otherwise, expressions such as $\partial V_{ij}^* / \partial T$ are always evaluated at type *ij*'s equilibrium insurance contract in the following.

Proposition 1.1. *Given the government's problem (1.6), the optimality conditions, assuming interior solutions, are*

$$(T): \quad 1 = \bar{b} + \sum_{i,j} n_{ij} g_{ij} \frac{\partial \beta_{ij}}{\partial T} \tag{1.10}$$

$$(\tau): \quad \frac{\tau}{1-\tau} = \frac{-Cov\,(wL, b) + \sum_{i,j} n_{ij} g_{ij} \left(\frac{\partial \beta_{ij}}{\partial \tau} + \frac{\partial \beta_{ij}}{\partial T} \bar{y} \right)}{\sum_{i,j} n_{ij} w_i L^*_{ij} \epsilon_{ij}} \tag{1.11}$$

$$(\alpha): \quad Cov \left(b, \frac{\partial d}{\partial \beta} \right) = -\sum_{i,j} n_{ij} g_{ij} \left(1 + \frac{\partial \beta_{ij}}{\partial \alpha} \right), \tag{1.12}$$

where $\bar{y} = \sum_{i,j} n_{ij} w_i L^*_{ij}$ *stands for average labor income,* $\bar{b} = \sum_{i,j} n_{ij} b_{ij}$ *for the population average of* b_{ij}*, and* $Cov(x, z) = \sum_{i,j} n_{ij}(x_{ij} - \overline{x})(z_{ij} - \overline{z})$ *for the covariance between two variables* x *and* z *in the population.*

Proof. See Appendix 1.A.3. □

Conditions (1.10)-(1.12) generalize the respective conditions that were obtained by Cremer and Pestieau (1996) for the case of perfect private insurance markets. They will be discussed in the next subsections.

1.3.3 Optimal Transfer T

We begin with considering condition (1.10). It differs from the condition that Cremer and Pestieau (1996) obtain only by the additional last term $\sum_{i,j} n_{ij} g_{ij}\, \partial \beta_{ij} / \partial T$. This term actually drops out if insurance markets are perfect. The individual contracts are then given by $\mathcal{I}_{ij} = (1 - \alpha, p_j(1 - \alpha)D)$, as everybody will buy full coverage for the remaining risk at an individually fair premium. These contracts are neither affected by T nor by τ. Condition (1.10) then states that the marginal gain from increasing the transfer, measured in terms of revenue, $\bar{b}$, should equal its cost.

In the case of imperfect insurance markets, increasing the transfer can have additional effects on the individuals' equilibrium contracts. For example, increasing the transfer may increase β_{ij} and d_{ij} for individuals with a positive marginal valuation of insurance. This would make a higher transfer T desirable.[18] To the contrary, if higher insurance coverage leads to large reductions in labor supply and therefore

[18]Statements of this type do of course only indicate tendencies of the additional effects due to market imperfections, since the optimal government policy is jointly determined by all optimality conditions.

$g_{ij} < 0$, imperfect insurance markets weaken the case for high transfers. This will be the case if individuals exhibit a large prudence.

1.3.4 *Optimal Tax Rate* τ

With the same argument as above, if neither T nor τ have an effect on the insurance market outcome, as for the case of perfect markets, condition (1.11) becomes the condition obtained by Cremer and Pestieau (1996).[19] It reflects the trade-off between efficiency and equity that is fundamental to the theory of optimal taxation. The numerator of the RHS reflects the goal of redistribution, since the population covariance between income and marginal social valuation can be interpreted as a welfare-based measure of inequality. A large negative covariance makes a higher tax rate more desirable. The denominator captures the distortionary effect of taxation. If labor supply reacts strongly to taxation (ϵ_{ij} large), optimal tax rates will be smaller. If taxation did not cause distortions, redistribution should take place until the correlation between income and marginal social valuation vanishes.

As before, the impact of taxation on the insurance market has to be taken into account in our more general setup. The term $\sum_{i,j} n_{ij} g_{ij} (\partial\beta_{ij}/\partial\tau + \bar{y}\partial\beta_{ij}/\partial T)$ captures these additional effects. If, for example, the tax rate has a desirable impact on the insurance market by increasing the insurance coverage for people who have a positive social marginal valuation of insurance, this effect has to be accounted for in favor of taxation. It is worth noting why this effect enters (1.11) as an on average compensated effect, i.e. why we find the term $\bar{y}\partial\beta_{ij}/\partial T$ in the brackets. The overall effect of τ on the market contracts is captured by $\partial\beta_{ij}/\partial\tau$. The revenue generated by a marginal increase in τ has the additional effect $\bar{y}\partial\beta_{ij}/\partial T$ in the insurance market and can be thought of as the "negative cost" of taxation.

1.3.5 *Optimal Social Insurance* α

Again, if insurance markets are perfect, condition (1.12) reduces to the respective condition in Cremer and Pestieau (1996), namely that $\text{Cov}(b, p) = 0$. Social insurance then simply crowds out private insurance ($\partial\beta_{ij}/\partial\alpha = -1$) so that the term on the RHS becomes zero. On the other hand, premiums are actuarially fair and (1.12) therefore states that the covariance between damage probability and

[19] Cremer and Pestieau (1996) derive the condition with only τ on the LHS but define the elasticity as $\epsilon_{ij} = w_i/L^*_{ij} \times \partial L^c_{ij}/\partial w^n_i$, such that their result is equivalent to ours.

marginal social valuation should be zero. This reflects that social insurance is a non-distortionary means of redistribution. High-risks benefit from social insurance since their private premiums would be larger than the social insurance contribution. The reverse holds for low-risks, so that increasing α redistributes from low to high risks and lowers the covariance between risk and marginal social valuation.[20] The government should do this until no correlation remains and the potential of social insurance for redistribution is exhausted.[21]

Condition (1.12) was derived under the assumption of an interior solution. It will become an inequality if the optimal α is a corner solution. Indeed, with perfect markets the optimal share of social insurance will always be one if high productivity individuals have lower damage probabilities (Cov$(p, w) < 0$), which is the empirically relevant case.[22] With full social insurance, individuals differ only with respect to their productivity, so that high risk types will still have the higher marginal social valuation due to their productivity disadvantage. Hence Cov$(b, p) > 0$ for all values of α and full social insurance is optimal. This leads us to the following corollary, which is the result found by Cremer and Pestieau (1996):

Corollary 1.1. *If Cov $(p, w) < 0$ and private insurance markets are perfect, the optimal social insurance level is $\alpha = 1$.*[23]

The model with perfect markets therefore requires a positive correlation between productivity and risk for less than full social insurance to be optimal. This case is empirically unappealing, so the model implies that partial social insurance is never optimal in real world situations.

Now consider the general version (condition (1.12)). Suppose first that private insurance premiums are still adjusted actuarially fairly if government policy changes the equilibrium coverage ($\partial d_{ij}/\partial \beta_{ij} = p_j D$). Even if the correlation between risk and social valuation is positive for all values of α, partial social insurance or even $\alpha = 0$ can now be optimal if the RHS of (1.12) is positive. In particular,

[20]This statement is true because of concavity of utility, which is the fundamental reason why redistribution is desirable from a utilitarian perspective.

[21]Note the analogy to the previous section where it was argued that a zero correlation between social valuation and a variable that the government can use to (indirectly) condition its policy on is optimal if the respective policy instrument causes no distortion.

[22]See Henriet and Rochet (2004) for some empirical evidence.

[23]In fact, the government would want to set $\alpha > 1$ under these circumstances. We do not consider this possibility due to strong moral hazard problems associated with overinsurance.

assume that social insurance indeed increases overall coverage for underinsured individuals ($\partial\beta_{ij}/\partial\alpha > -1$). The often discussed *efficiency effect* of mitigating underinsurance is then present. It is captured by the first, positive term in g_{ij} (see equation (1.8)): underinsured individuals experience an increase in utility if their overall coverage grows and premiums are adjusted fairly at the margin. However, there is a second, negative term in g_{ij} which captures the *precautionary labor effect*: individuals will react to the reduction of risk by reducing labor supply, as shown in Lemma 1.1. Since this reduces the revenues from income taxation, it will negatively affect social welfare. When this effect dominates, the sign of g_{ij} is negative and the RHS of (1.12) becomes positive. The argument for social insurance is then weakened by adverse selection compared to situations with perfect insurance markets.

Finally, the fact that the adjustment of the private premium to changes in coverage matters for the covariance in (1.12) points at the logic of redistribution via social insurance. It is not damage probability per se but the possible savings on private premiums that matter for redistribution. While both are the same if markets are perfect, risk and premium can diverge under adverse selection and substantially change the *redistribution effect* of social insurance. If, for example, all individuals pay the same premium in a pooling contract, the covariance in (1.12) is zero. In this case, there is no justification for social insurance from the point of view of redistribution.

1.4 Imperfect Insurance Markets

While the dependence of the equilibrium insurance contracts on the policy parameters has been left unrestricted so far, we now proceed to show how such relations emerge from endogenizing the insurance market equilibrium.

We assume that insurance companies have no information on individual risks p_j but can observe the individual productivity levels w_i. This assumption is to keep our exposition as simple as possible since otherwise we would have to deal with a problem of two-dimensional adverse selection.[24] It allows us to divide the private

[24]The same assumption is taken by Boadway et al. (2006). We are aware that it still entails an informational inconsistency, since government and insurers do not have the same set of information. A similar inconsistency already underlies the classical income tax model by Mirrlees (1971), where firms pay wages according to individual productivities whereas the government can only observe total income. With our information structure, insurance companies have the same information as the firms on the labor market, so that any asymmetry within the private sector is eliminated.

insurance market into W sub-markets, one for each productivity level. Adverse selection in each of those markets can be modeled using a variety of game theoretic approaches. In the following, we demonstrate how the general optimality condition for social insurance can be applied to the equilibrium concepts developed by Rothschild and Stiglitz (1976), Wilson (1977) and Miyazaki (1977).[25]

Before we proceed, a remark about how to include insurance markets into our model is necessary. Since labor supply is endogenous, it reacts to both changes in expected consumption and changes in risk as shown in Section 1.2. This means that we have to account for variations in labor supply as we move through the insurance contract space, for example to determine the shape and crossing properties of indifference curves. A more detailed discussion of this topic is relegated to Appendix 1.A.4. There we derive sufficient conditions under which the insurance markets work as in the standard case without endogenous labor supply.[26] This will be assumed for the rest of the chapter.

1.4.1 Rothschild-Stiglitz Separating Equilibria

We first consider separating equilibria as suggested by Rothschild and Stiglitz (1976) in each of the W private insurance markets, starting in a situation without social insurance. Note that when the separating equilibrium exists for $\alpha = 0$, a separating equilibrium of the same type will also exist for all other values of $\alpha \in (0, 1)$.[27]

High risks obtain full coverage at an individually fair premium, i.e. $\beta_{iH} = 1 - \alpha$ and $d_{iH} = p_H \beta_{iH} D$. The low risks' equilibrium contract lies on the low risks' zero-profit line and is such that the high risks' incentive compatibility constraint is just binding. Formally, β_{iL} solves

$$V_{iH}^*(\tau, T, \alpha, 1 - \alpha, p_H(1 - \alpha)D) = V_{iH}^*(\tau, T, \alpha, \beta_{iL}, p_L \beta_{iL} D), \tag{1.13}$$

[25] A similar analysis (based on their different timing assumption for labor supply mentioned above) has been performed by Boadway et al. (2006), who only consider Rothschild-Stiglitz equilibria. By applying the results of Proposition 1.1 to different models of adverse selection, we try to provide a more general analysis and identify the robust effects.

[26] In particular, conditions are derived for indifference curves to be concave and single-crossing.

[27] See for example Eckstein et al. (1985). This property ensures that the government's problem is actually continuous. Existence of the Rothschild-Stiglitz equilibrium can be guaranteed by applying the equilibrium concept developed by Riley (1979).

which states that high risks cannot do better by choosing the contract $(\beta_{iL}, p_L\beta_{iL}D)$ instead of their own contract. Utility from mimicking is given by the RHS of (1.13) and shall be denoted with the shortcut $\tilde{V}_{iH}$ in the following. Clearly, the low risks are underinsured, i.e. $\beta_{iL} < 1 - \alpha$.

We first derive the net social marginal valuation of insurance, g_{ij}, for the different types. As all individuals pay an actuarially fair premium, the premium adjusted effect of insurance coverage on labor supply reduces to the pure precautionary effect

$$\left.\frac{\partial L^*_{ij}}{\partial \beta}\right|_A = -\frac{\partial L^*_{ij}}{\partial T}\frac{\partial \Psi_{ij}}{\partial \beta} < 0. \tag{1.14}$$

In addition, (1.14) completely vanishes for the high risks because $\partial\Psi_{ij}/\partial\beta = 0$ at full insurance. Next, we need to examine

$$\left.\frac{\partial V^*_{ij}}{\partial \beta}\right|_A = \frac{\partial V^*_{ij}}{\partial \beta} - p_j D\frac{\partial V^*_{ij}}{\partial T}.$$

This is again zero for the high risks, a direct implication of the fact that they obtain their optimal fair contract. The low risks, however, derive positive utility from an increase in coverage with fair adjustment of the premium. This follows immediately from risk aversion and the fact that they are underinsured. Hence we have the following net social marginal valuation of insurance for the two risk types:

$$g_{iL} = \frac{1}{\gamma}\left(\frac{\partial V^*_{iL}}{\partial \beta} - p_L D\frac{\partial V^*_{iL}}{\partial T}\right) - \tau w_i \frac{\partial L^*_{iL}}{\partial T}\frac{\partial \Psi_{iL}}{\partial \beta} \quad \text{and} \quad g_{iH} = 0. \tag{1.15}$$

It is zero for the high risks as they obtain their first-best contract. For the low-risks types, the two counteracting welfare effects of variations in insurance coverage discussed in section 1.3.5 are present. They benefit from additional coverage at a fair premium, but at the same time supply less labor due to reduced risk. This reduces tax revenue and therefore welfare. The overall sign of g_{iL} depends notably on the size of the coefficient of prudence.

To determine the optimal amount of social insurance, we need to know how it affects individuals' *overall* coverage. Clearly, as the high risks are always fully insured, social insurance simply crowds out their private insurance: $\partial\beta_{iH}/\partial\alpha = -1$.

For the low risks, implicit differentiation of (1.13) gives after some rearrangements[28]

$$\frac{\partial \beta_{iL}}{\partial \alpha} = -\frac{\partial \tilde{V}_{iH}/\partial \beta - \left((p_H - \bar{p})\partial V^*_{iH}/\partial T + \bar{p}\,\partial \tilde{V}_{iH}/\partial T \right) D}{\partial \tilde{V}_{iH}/\partial \beta - p_L D \partial \tilde{V}_{iH}/\partial T}. \tag{1.16}$$

The denominator of (1.16) is positive, because the single-crossing property and concavity of indifference curves together imply

$$\frac{\partial \tilde{V}_{iH}/\partial \beta}{\partial \tilde{V}_{iH}/T} > \frac{\partial V^*_{iL}/\partial \beta}{\partial V^*_{iL}/T} > p_L D.$$

In addition, due to

$$(p_H - \bar{p})\frac{\partial V^*_{iH}}{\partial T} + \bar{p}\frac{\partial \tilde{V}_{iH}}{\partial T} > p_L \frac{\partial \tilde{V}_{iH}}{\partial T},$$

the numerator is smaller than the denominator in (1.16), which establishes

$$\frac{\partial \beta_{iL}}{\partial \alpha} > -1. \tag{1.17}$$

An increase in social insurance unambiguously increases overall coverage for low risks. We summarize our results in the following Corollary of Proposition 1.1:

Corollary 1.2. *With Rothschild-Stiglitz separating equilibria on each of the W private insurance markets, the optimality condition (1.12) becomes*

$$D\,Cov\,(b,p) = -\sum_i n_{iL} g_{iL} \left(1 + \frac{\partial \beta_{iL}}{\partial \alpha} \right), \tag{1.18}$$

where g_{iL} is given by (1.15) and $1 + \partial \beta_{iL}/\partial \alpha > 0$.

Indeed, social insurance redistributes as in the case of perfect markets and additionally has positive effects by reducing underinsurance. The negative effect of reducing labor supply is present as well, however. Even for Cov $(p, w) < 0$, the optimal social insurance level α can be less than one if households are sufficiently prudent. The precautionary labor effect alone can justify interior levels of social insurance or even complete renunciation of a social insurance system.

[28]The derivation makes use of the fact that effects of α can be restated as appropriately compensated effects of β as shown in Appendix 1.A.3.

1.4.2 Wilson Pooling Equilibria

To what extent do the results in the previous subsection 1.4.1 depend on the particularities of the Rothschild-Stiglitz equilibrium? In order to answer this question, we now consider an alternative concept going back to Wilson (1977) and examine the case of pooling equilibria.[29]

The problem with Wilson pooling equilibria is that after the introduction of social insurance, separating equilibria might emerge on some of the W private markets even if there has been pooling in the absence of social insurance. We therefore would have to account for the fact that, as we gradually increase α, some markets experience a discrete jump into a separating equilibrium earlier than the others. A mix of marginal conditions based on pooling and separating equilibria would then need to be considered. To be able to highlight the effects of pooling in contrast to the previous subsection, we do not analyze this complication but assume that the optimal solution entails pooling in all W markets.[30] Also, the results are useful for understanding the next section, which deals with an equilibrium concept that combines separating and pooling components.

For each productivity group i, the Wilson pooling equilibrium is such that both risk types choose the low risks' preferred contract on the zero profit line for the whole group. Consequently, it is determined by

$$\beta_i = \arg\max_{\delta} V_{iL}^*(\tau, T, \alpha, \delta, \bar{p}_i \delta D). \tag{1.19}$$

There are two fundamental differences to the Rothschild-Stiglitz separating equilibrium. First, the whole population is underinsured. Second, no individual pays an individually fair premium, but cross-subsidization from low to high risks occurs.

[29] As was shown by Wilson (1977), if a Rothschild-Stiglitz separating equilibrium exists, it is also an equilibrium in Wilson's sense. Otherwise, a pooling equilibrium exists where both risk types choose the same contract.

[30] For a single insurance market, Eckstein et al. (1985) have shown that a Wilson pooling equilibrium is never efficient. A Pareto improvement can be generated using social insurance, which in turn induces a separating equilibrium on the remaining market. One may therefore wonder why optimality conditions should ever be based on pooling equilibria. There are two reasons. First, our model includes W markets with generally *different* pooling equilibria. Since social insurance cannot be differentiated according to productivities, such equilibria can be Pareto efficient in our framework. Second, precautionary motives can make social insurance harmful and render pooling efficient even in a single market.

Turning first to the effect of a premium-adjusted change in β_i on labor supply, we obtain

$$\left.\frac{\partial L_{ij}^*}{\partial \beta}\right|_A = \left((p_j - \bar{p}_i)D - \frac{\partial \Psi_{ij}}{\partial \beta}\right)\frac{\partial L_{ij}^*}{\partial T}. \tag{1.20}$$

Compared to (1.14), there is an additional income effect from the unfair premium. High risks unambiguously reduce their labor supply, since both the better than fair premium and the precautionary effect work in the same direction. The labor effect is in general undetermined for the low risks but also negative if they are sufficiently prudent. Similarly, individuals' utility is affected as follows:

$$\left.\frac{\partial V_{ij}^*}{\partial \beta}\right|_A = \frac{\partial V_{ij}^*}{\partial \beta} - \bar{p}_i D \frac{\partial V_{ij}^*}{\partial T}.$$

This expression for the low risks corresponds to the first order condition of (1.19) and is zero. Single-crossing implies that it is positive for the high risks. Thus, in a reversal of our findings based on the separating equilibrium, only the high risks derive a utility gain from increased coverage. With these insights, we can again summarize the net social marginal valuation of insurance for the two types. For the low risks, we are left with

$$g_{iL} = \tau w_i \left((p_L - \bar{p}_i)D - \frac{\partial \Psi_{iL}}{\partial \beta}\right)\frac{\partial L_{iL}^*}{\partial T}, \tag{1.21}$$

which only contains the ambiguous labor supply distortion explained above. The high risks are characterized by

$$g_{iH} = \frac{1}{\gamma}\left(\frac{\partial V_{iH}^*}{\partial \beta} - \bar{p}_i D \frac{\partial V_{iH}^*}{\partial T}\right) + \tau w_i \left((p_H - \bar{p}_i)D - \frac{\partial \Psi_{iH}}{\partial \beta}\right)\frac{\partial L_{iH}^*}{\partial T}. \tag{1.22}$$

Their net social marginal gain from insurance captures the two opposite effects on utility and labor supply and is ambiguous as well. Yet, both terms crucially depend on the strength of the precautionary motive.

Under the Wilson pooling equilibrium, the effect of α on private coverage cannot be signed without more specific assumptions on utility functions. Because empirical studies suggest that social insurance does not completely crowd out private insurance, it seems sensible to focus on the case in which again $\partial \beta_i / \partial \alpha > -1$ holds.[31]

[31] See, for example, the work of Cutler and Gruber (1996) who find a crowding-out rate of roughly 50%.

Corollary 1.3. *With Wilson pooling equilibria on each of the W private insurance markets, the optimality condition (1.12) becomes*

$$D\,Cov\,(b, \bar{p}_i) = -\sum_{i,j} n_{ij} g_{ij} \left(1 + \frac{\partial \beta_i}{\partial \alpha}\right), \tag{1.23}$$

where g_{iL} is given by (1.21) and g_{iH} by (1.22).

The crucial difference to (1.18) relates to the redistributive impact. Instead of $\mathrm{Cov}(b, p)$, we find the covariance between b and the *average* damage probabilities for each productivity group. Clearly, social insurance is no longer able to directly redistribute between risks, as this is already achieved by cross-subsidization in the market. The only remaining redistributive effect is across productivity classes if they are characterized by different private insurance premia. The reduced redistributive potential and the already occurring equalization of types in the market gives rise to a smaller covariance than before and considerably weakens the case for social insurance.

The results lead us to the conclusion that both the redistribution effect and the precautionary labor effect work in the same direction and may even make the corner solution without social insurance optimal.

1.4.3 Miyazaki-Wilson Equilibria

Finally, we examine an equilibrium concept that goes back to Wilson (1977) and Miyazaki (1977).[32] It has two interesting properties. First, as was shown by Crocker and Snow (1985), it is always second-best efficient. Second, it is a separating equilibrium which still involves cross-subsidization from low to high risks. Its analysis is therefore essentially a careful combination of the results for the Rothschild-Stiglitz separating and the Wilson pooling equilibria.

As in Section 1.4.1, high risks will obtain full coverage ($\beta_{iH} = 1 - \alpha$). Their premium can, however, contain a subsidy denoted by y_i, so that $d_{iH} = p_H \beta_{iH} D - y_i$. This subsidy is paid by the low risks, so that their premium for any given coverage δ becomes $p_L \delta D + (n_{iH}/n_{iL}) y_i$. Their contract must then fulfill the incentive compatibility constraint

$$V_{iH}^*(\tau, T, \alpha, 1-\alpha, p_H(1-\alpha)D - y_i) = V_{iH}^*(\tau, T, \alpha, \delta, p_L \delta D + (n_{iH}/n_{iL}) y_i), \tag{1.24}$$

[32] It is often also associated with Spence (1978).

where the RHS will again be denoted by $\tilde{V}_{iH}$. Equation (1.24) therefore defines the subsidy $y_i(\delta)$ that has to be paid for δ to be incentive feasible for the low risks. To see how δ and y_i relate, implicitly differentiate (1.24) to obtain

$$\frac{\partial y_i}{\partial \delta} = \frac{\partial \tilde{V}_{iH}/\partial \beta - p_L D \, \partial \tilde{V}_{iH}/\partial T}{\partial V^*_{iH}/\partial T + (n_{iH}/n_{iL})\partial \tilde{V}_{iH}/\partial T} > 0. \tag{1.25}$$

Increasing the subsidy slackens the incentive constraint and makes it possible to increase the low risks' coverage. Equilibrium coverage is then determined as

$$\begin{aligned} \beta_{iL} = \arg\max_{\delta} V_{iL}(\tau, T, \alpha, \delta, p_L \delta D + (n_{iH}/n_{iL}) y_i(\delta)) \\ \text{s. t. } \; y_i(\delta) \geq 0. \end{aligned} \tag{1.26}$$

If the constraint is binding, the Miyazaki-Wilson equilibrium is identical to the Rothschild-Stiglitz separating equilibrium and our results from section 1.4.1 apply unaltered. Otherwise, the equilibrium differs by the subsidy $y_i > 0$. If the policy parameters change the low risks' coverage β_{iL}, the corresponding changes in y_i have to be taken into account in addition to what has been derived in section 1.4.1.

Note first that this makes a slight modification of the optimality condition (1.12) necessary. In deriving Proposition 1.1 it has been assumed that no interdependencies of this kind exist, i.e. that each type's premium does only depend on its own coverage. If the derivations in Appendix 1.A.3 are repeated allowing for this interrelation, it turns out that the modified optimality condition becomes

$$\text{Cov}\left(b, \frac{\partial d}{\partial \beta}\right) = -\sum_{i,j} n_{ij} g_{ij} \left(1 + \frac{\partial \beta_{ij}}{\partial \alpha}\right) - \sum_{i} n_{iH} b_{iH} \left(\frac{\partial y_i}{\partial \delta} \frac{\partial \beta_{iL}}{\partial \alpha}\right). \tag{1.27}$$

This condition captures the fact that whenever β_{iL} increases, the high risks profit from a larger subsidy, which has to be weighted by their social valuation of income.

The rest of the analysis proceeds as before. As for the high risks' net social marginal valuation of insurance, we again find $g_{iH} = 0$, with the same argument as in Section 1.4.1. Low risks' utility is not directly affected by a small premium-adjusted increase in β_{iL}. The first order condition of problem (1.26) implies that this effect is zero in the considered case in which the constraint is slack. This confirms that social insurance can have no positive efficiency impact in case the market equilibrium is second-best. Yet, the standard distortions of labor supply still remain present. The premium change triggered by an increase in β_{iL} can be decomposed into a marginally fair adjustment and the unfair increase of the subsidy. The latter is not compensated by a lower expected damage. It counteracts the precautionary effect but will not dominate if the low risks' prudence is large. We summarize

$$g_{iL} = \tau w_i \frac{\partial L_{iL}^*}{\partial T}\left[-\frac{\partial \Psi_{iL}}{\partial \beta} - \frac{n_{iH}}{n_{iL}}\frac{\partial y_i}{\partial \delta}\right]. \tag{1.28}$$

As before, signing $1 + \partial\beta_{iL}/\partial\alpha$ is not possible in general. It again seems reasonable to consider the case with less than complete crowding out, so that $-1 < \partial\beta_{iL}/\partial\alpha < 0$.

Corollary 1.4. *With Miyazaki-Wilson equilibria on each of the W private insurance markets, the optimality condition for α is*

$$Cov\left(b, \frac{\partial d}{\partial \beta}\right) = -\sum_i \left[n_{iL} g_{iL}\left(1 + \frac{\partial \beta_{iL}}{\partial \alpha}\right) + n_{iH} b_{iH}\left(\frac{\partial y_i}{\partial \delta}\frac{\partial \beta_{iL}}{\partial \alpha}\right)\right] \tag{1.29}$$

where g_{iL} is given by (1.28) and $\partial y_i/\partial\delta$ by (1.25).

Two major differences to Section 1.4.1 have to be accounted for when interpreting (1.29). First, g_{iL} does no longer include a positive utility effect on the low risks from higher coverage, but only the labor supply distortions. Second, the redistributive impact of social insurance is reduced due to the "pooling component" in the equilibrium. The individuals' marginal savings on private insurance as captured by $\partial d/\partial\beta$ in the covariance term in (1.29) contain the actuarially fair part $p_j D$ as in the Rothschild-Stiglitz case. For the low risks, however, the additional term $\partial y_i/\partial\delta > 0$ appears. If social insurance (partially) replaces private insurance, the subsidy that they already paid to the high risks in the market drops and "private redistribution" is crowded out. The effect on the beneficiaries of the subsidy is then captured by the additional positive term on the RHS of (1.29). In this respect, the present equilibrium indeed represents an intermediate case between the simple separating and the complete pooling scenarios. The fact that no positive efficiency effects but only labor supply distortions arise and the redistributive power is reduced lets us conclude that a lower level of social insurance compared to a situation with (first best) efficient insurance markets will again be optimal if precautionary motives are sufficiently strong.

1.5 Conclusions

We have developed a theory of optimal taxation and social insurance in the presence of imperfect private insurance markets. While the problem of taxation requires to model the households' choice of labor supply endogenously, inefficient insurance markets imply that they have to take this decision under risk. Hence, a theory of labor supply under uncertainty provides the basis for our analysis. As we have shown, there exists a motive for precautionary labor under general and meaningful circumstances.

The integration of imperfect insurance markets into a model of taxation and social insurance allowed us to show how the optimality conditions for public policy based on efficient insurance markets have to be modified. Notably, the strength of the precautionary labor motive turns out to be crucial in determining whether social insurance should be higher or lower compared to earlier results. Social insurance might have efficiency-enhancing effects by reducing underinsurance. At the same time, larger overall insurance coverage leads individuals to reduce their labor supply, which emerges as an important repercussion. Furthermore, social insurance suffers from substantially reduced redistributive power when equilibria with cross-subsidization prevail on markets. Finally, even the positive efficiency effects vanish in a second-best market equilibrium. Using specific equilibrium concepts for the insurance markets, we illustrated these results and showed that it may even be optimal to completely renounce on social insurance as a policy device and only use income taxation to achieve redistributive objectives. This is in stark contrast to conjectures of a previous literature where social insurance provided a means of distortion-free redistribution.

Our contribution has raised issues for further research. Our theory of precautionary labor may provide an important and so far unexplored tool for analyzing a variety of other economic problems such as optimal labor contracts, unemployment and macroeconomic fluctuations. In addition, the interaction between precautionary labor and precautionary savings might lead to new and interesting effects in dynamic models that account for the interaction of labor, insurance, and capital markets.

1.A Appendix

1.A.1 Risk Aversion and Prudence

The Arrow-Pratt coefficient of absolute risk aversion for $u(c)$ is defined as

$$r_A(c) = -\frac{u''(c)}{u'(c)} > 0. \tag{1.30}$$

Absolute risk aversion is therefore constant or decreasing if

$$\frac{\partial r_A(c)}{\partial c} = -\frac{u'(c)u'''(c) - [u''(c)]^2}{[u'(c)]^2} \leq 0. \tag{1.31}$$

A necessary condition for this to be satisfied is that the numerator of (1.31) is non-negative. For this in turn it is necessary that $u'''(c)$ is positive. Analogously, the coefficient of relative risk aversion is given by

$$r_R(c) = -\frac{c\,u''(c)}{u'(c)} > 0. \tag{1.32}$$

Relative risk aversion is constant or decreasing if

$$\frac{\partial r_R(c)}{\partial c} = -\frac{u'(c)[c\,u'''(c) + u''(c)] - c\,[u''(c)]^2}{[u'(c)]^2} \leq 0. \tag{1.33}$$

With the same argument as above, this can only be satisfied if the numerator is positive, which requires $u'''(c)$ to be positive.

1.A.2 Proof of Lemma 1.2

Consider a damage D that occurs with probability p, and of which a share β is insured (free of charge). This defines a Bernoulli random variable with expectation $p\,(1-\beta)D$ and variance $p\,(1-p)[(1-\beta)D]^2$. The first order condition for optimal labor supply L^* is

$$w[p\,u'(wL^* + T - (1-\beta)D) + (1-p)\,u'(wL^* + T)] = -v'(L^*). \tag{1.34}$$

The income effect can be derived by implicitly differentiating (1.34):

$$\frac{\partial L^*}{\partial T} = -\frac{w\,[pu''(wL^* + T - (1-\beta)D) + (1-p)u''(wL^* + T)]}{SOC}, \tag{1.35}$$

where SOC stands for the second derivative of the objective with respect to L and is negative. Therefore, the income effect is negative, which implies that leisure is a normal good.[33] Implicit differentiation of (1.34) w.r.t. β yields after some rearrangements

$$\frac{\partial L^*}{\partial \beta} = \frac{\partial L^*}{\partial T} pD - \frac{w[u''(wL^* + T - (1-\beta)D) - u''(wL^* + T)]}{SOC} p(1-p)D. \quad (1.36)$$

The income effect due to decreased expected damage is already visible as the first term on the RHS of (1.36). Substituting $\Delta u''(\beta)$ for $u''(wL^* + T) - u''(wL^* + T - (1-\beta)D)$, $\mathrm{E}[u''(.)]$ for $p\, u''(wL^* + T - (1-\beta)D) + (1-p)\, u''(wL^* + T)$, together with the income effect (1.35) and the first derivative of the variance w.r.t to β, (1.36) can be transformed to

$$\frac{\partial L^*}{\partial \beta} = \frac{\partial L^*}{\partial T} \left[pD - \frac{\partial \Psi}{\partial \beta} \right], \quad (1.37)$$

where

$$\frac{\partial \Psi}{\partial \beta} = \left(-\frac{\Delta u''(\beta)/(1-\beta)D}{\mathrm{E}[u''(.)]} \right) \left(\frac{1}{2} \frac{\partial \mathrm{Var}}{\partial \beta} \right). \quad (1.38)$$

1.A.3 Proof of Proposition 1.1

Assuming interior solutions, the three first order conditions for problem (1.6) are

$$(T): \quad \sum_{i,j} n_{ij} \frac{\partial V_{ij}^{**}}{\partial T} + \gamma \sum_{i,j} n_{ij} \left(-1 + \tau w_i \frac{\partial L_{ij}^{**}}{\partial T} \right) = 0 \quad (1.39)$$

$$(\tau): \quad \sum_{i,j} n_{ij} \frac{\partial V_{ij}^{**}}{\partial \tau} + \gamma \sum_{i,j} n_{ij} \left(w_i L_{ij}^{**} + \tau w_i \frac{\partial L_{ij}^{**}}{\partial \tau} \right) = 0 \quad (1.40)$$

$$(\alpha): \quad \sum_{i,j} n_{ij} \frac{\partial V_{ij}^{**}}{\partial \alpha} + \gamma \sum_{i,j} n_{ij}\, \tau w_i \frac{\partial L_{ij}^{**}}{\partial \alpha} = 0, \quad (1.41)$$

where γ is the Lagrange multiplier associated with the revenue constraint.

Consider first the condition for the optimal transfer T. The effect of T on L_{ij}^{**} can be reduced to effects on L_{ij}^{*} by explicitly taking into account its effect on the private insurance market equilibrium:

$$\frac{\partial L_{ij}^{**}}{\partial T} = \frac{\partial L_{ij}^{*}}{\partial T} + \left. \frac{\partial L_{ij}^{*}}{\partial \beta} \right|_A \frac{\partial \beta_{ij}}{\partial T}. \quad (1.42)$$

The same decomposition can be applied to V_{ij}^{**} and it is also applicable to the effects of the other government parameters. Using this, as well as the definition of the two concepts of marginal social valuation (b_{ij} and g_{ij}) from section 3, the optimality condition (1.10) follows after some simple rearrangements.

[33]This results from the fact that utility is separable.

Consider next the first order condition for the optimal level of social insurance α. Again, the effects of α on L_{ij}^{**} and V_{ij}^{**} can be reduced to effects on L_{ij}^{*} and V_{ij}^{*}. After some rearrangements, one obtains

$$\sum_{i,j} n_{ij} \left(\frac{1}{\gamma} \frac{\partial V_{ij}^*}{\partial \alpha} + \tau w_i \frac{\partial L_{ij}^*}{\partial \alpha} \right) + \sum_{i,j} n_{ij} g_{ij} \frac{\partial \beta_{ij}}{\partial \alpha} = 0. \tag{1.43}$$

The first term on the LHS of (1.43) can further be transformed by noting that effects of α on labor supply can be expressed as effects of β as follows:

$$\frac{\partial L_{ij}^*}{\partial \alpha} = \left. \frac{\partial L_{ij}^*}{\partial \beta} \right|_A + (d'_{ij} - \bar{p}D) \frac{\partial L_{ij}^*}{\partial T}, \tag{1.44}$$

where d'_{ij} stands short for $\partial d_{ij}/\partial \beta_{ij}$. Equation (1.44) follows from the fact the changes in social and private insurance differ only with respect to their different premiums. The same decomposition holds for indirect utility. Substituting this into (1.43) yields

$$\sum_{i,j} n_{ij} g_{ij} \left(1 + \frac{\partial \beta_{ij}}{\partial \alpha} \right) + \sum_{i,j} n_{ij} b_{ij} (d'_{ij} - \bar{p}D) = 0. \tag{1.45}$$

After adding and subtracting $\sum_{i,j} n_{ij} \bar{b} (d'_{ij} - \bar{p}D)$ one obtains

$$\sum_{i,j} n_{ij} g_{ij} \left(1 + \frac{\partial \beta_{ij}}{\partial \alpha} \right) + \sum_{i,j} n_{ij} (b_{ij} - \bar{b})(d'_{ij} - \bar{p}D) + \bar{b} \sum_{i,j} n_{ij} (d'_{ij} - \bar{p}D) = 0. \tag{1.46}$$

The last term on the LHS of (1.46) is equal to zero since aggregate profits of insurance companies in a competitive market equilibrium are zero. Therefore, higher insurance coverage for all individuals will be accompanied by adjustments in the premiums such that additional revenues equal additional expected insurance payments on the population average.[34]

Finally, the first order condition for the optimal tax rate τ is to be examined. Following the same first steps as before yields

$$\sum_{i,j} n_{ij} g_{ij} \frac{\partial \beta_{ij}}{\partial \tau} + \sum_{i,j} n_{ij} \left(\frac{1}{\gamma} \frac{\partial V_{ij}^*}{\partial \tau} + \tau w_i \frac{\partial L_{ij}^*}{\partial \tau} \right) + \sum_{i,j} n_{ij} w_i L_{ij}^* = 0. \tag{1.47}$$

[34] This holds under a vast majority of insurance market equilibria, among them all those discussed in section 1.4.

Since labor supply is chosen optimally, the effect of τ on V_{ij}^* boils down to a pure income effect. The tax rate effect on labor supply can be transformed using the Slutzky decomposition (1.9). After some rearrangements one obtains

$$\frac{\tau}{1-\tau} = \frac{\sum_{i,j} n_{ij} g_{ij} \frac{\partial \beta_{ij}}{\partial \tau} + \sum_{i,j} n_{ij} w_i L_{ij}^* - \sum_{i,j} n_{ij} b_{ij} w_i L_{ij}^*}{\sum_{i,j} n_{ij} w_i L_{ij}^* \epsilon_{ij}}. \tag{1.48}$$

According to the first order condition (1.10) for the optimal transfer T, the term $\bar{b} + \sum_{i,j} n_{ij} g_{ij}\, \partial \beta_{ij}/\partial T$ equals one and can therefore simply be multiplied with the second sum in the numerator of (1.48). After some rearrangements, the optimality condition (1.11) follows.

1.A.4 Labor Supply and Demand for Insurance

The first step to understanding individuals' preferences for insurance is to reexamine how labor supply reacts to variations in the insurance contract. As has been shown in Section 1.2, actuarially unfair changes in the insurance contract, which correspond to changes in expected income, trigger standard income effects. In addition, changes in variance cause precautionary effects. Thus, when considering an individual's preferences for insurance, we need to account for changes in labor supply and thus consumption levels as we move along an indifference curve in the (β, d)-space. The endogeneity of labor supply may alter the shape of indifference curves compared to the canonical models.

At a contract (β, d), consumption in case of loss is $c_{ij}^0 = w_i L_{ij}^* - (1-\beta)D - d$ and $c_{ij}^1 = w_i L_{ij}^* - d$ otherwise. Consider the slope of an indifference curve of an individual with productivity w_i and risk p_j in this contract

$$\left.\frac{\partial d}{\partial \beta}\right|_{V_{ij}^*=\bar{V}} = \mathrm{MRS}_{ij}(\beta, d) = \frac{D p_j u'(c_{ij}^0)}{p_j u'(c_{ij}^0) + (1-p_j) u'(c_{ij}^1)} > 0, \tag{1.49}$$

which is positive as in the standard model.[35] Note also that $\mathrm{MRS}_{ij} = p_j D$ at any full coverage contract (where $c_{ij}^0 = c_{ij}^1$), an additional result that carries over from the standard model. Next, we examine concavity of indifference curves in the (β, d)-space. In order to see how (1.49) changes as we move up on an indifference curve

[35] Clearly, indifference curves are still continuous and differentiable.

$d_{ij}(\beta)$, we need to evaluate the sign of

$$\frac{\partial \text{MRS}_{ij}(\beta, d_{ij}(\beta))}{\partial \beta} = p_j(1-p_j)D\frac{u''(c^0_{ij})u'(c^1_{ij})\frac{\partial c^0_{ij}}{\partial \beta} - u''(c^1_{ij})u'(c^0_{ij})\frac{\partial c^1_{ij}}{\partial \beta}}{(\partial V^*_{ij}/\partial d)^2}. \quad (1.50)$$

Since both $u''(c^0_{ij})u'(c^1_{ij})$ and $u''(c^1_{ij})u'(c^0_{ij})$ are negative, the expression (1.50) is negative whenever $\partial c^1_{ij}/\partial\beta < 0 < \partial c^0_{ij}/\partial\beta$, implying concavity of the indifference curves. This condition is clearly satisfied if labor supply is fixed exogenously, because the fundamental purpose of insurance is to redistribute from the state without damage to the state where the loss occurs. In the present model with endogenous income, it puts an upper bound on both the precautionary and the income effect. The income effect in response to an (unfair) increase in the insurance contract along an indifference curve is not allowed to be strong enough to let c^0_{ij} grow.[36] On the other hand, the precautionary effect of a reduced risk is not allowed to be strong enough to let c^1_{ij} shrink.

A second condition that needs to be satisfied to make adverse selection problems tractable is the single-crossing property. It implies that (within each productivity group) the high risk type is characterized by a steeper indifference curve everywhere in the (β, d)-space. Formally, the marginal rate of substitution must be increasing in p_j. As can be seen from (1.49), this would clearly hold if labor supply did not depend on p_j. However, by implicitly differentiating the first order condition (1.1) of the individual labor choice problem, we obtain for any $\beta < 1$

$$\frac{\partial L^*_{ij}}{\partial p_j} = -\frac{(1-\tau)w_i\left(u'(c^0_{ij}) - u'(c^1_{ij})\right)}{SOC} > 0 \quad (1.51)$$

because of risk aversion and the second order condition $SOC < 0$. Hence at any given contract with less than full coverage, high risk individuals supply more labor than low risks, so that $c^0_{iH} > c^0_{iL}$ and $c^1_{iH} > c^1_{iL}$. If preferences exhibit increasing risk aversion, this amplifies the effect of damage probability on the marginal willingness to pay for insurance, while there is no effect of labor supply under CARA.[37]

[36]The assumption of decreasing absolute risk-aversion would be a sufficient condition for this to hold.

[37]These statements are easily proven by considering the effect of increasing labor supply on (1.49).

In both cases, single-crossing clearly holds. If risk aversion is decreasing, adjustments of labor supply generate an opposite effect on the marginal rate of substitution (1.49). It can easily be shown, however, that single-crossing is still satisfied when p_H is sufficiently large while p_L is sufficiently small. This follows because the marginal rate of substitution converges to D as $p \to 1$ while it converges to 0 as $p \to 0$, irrespective of the endogeneity of labor supply. Throughout section 1.4 we assume both concavity and single-crossing.

References of Chapter 1

ABEL, A. (1988): "The Implications of Insurance for the Efficacy of Fiscal Policy," *Journal of Risk and Insurance*, 55, 339–378.

BOADWAY, R., M. LEITE-MONTEIRO, M. MARCHAND, AND P. PESTIEAU (2006): "Social Insurance and Redistribution with Moral Hazard and Adverse Selection," *Scandinavian Journal of Economics*, 108, 279–298.

CREMER, H. AND P. PESTIEAU (1996): "Redistributive Taxation and Social Insurance," *International Tax and Public Finance*, 3, 281–295.

CROCKER, K. AND A. SNOW (1985): "The Efficiency of Competitive Equilibria in Insurance Markets with Asymmetric Information," *Journal of Public Economics*, 26, 207–219.

CUTLER, D. AND J. GRUBER (1996): "Does Public Insurance Crowd Out Private Insurance?" *Quarterly Journal of Economics*, 111, 391–430.

EATON, J. AND H. S. ROSEN (1979): "Notes on Optimal Wage Taxation and Uncertainty," *NBER Working Paper No. 388*.

——— (1980): "Labor Supply, Uncertainty and Efficient Taxation," *Journal of Public Economics*, 14, 365–374.

ECKSTEIN, Z., M. EICHENBAUM, AND D. PELED (1985): "Uncertain Lifetimes and the Welfare Enhancing Properties of Annuity Markets and Social Security," *Journal of Public Economics*, 26, 303–326.

FLODEN, M. (2006): "Labor Supply and Saving under Uncertainty," *Economic Journal*, 116, 721–737.

HARTWICK, J. M. (2000): "Labor Supply under Wage Uncertainty," *Economics Letters*, 68, 319–325.

HENRIET, D. AND J.-C. ROCHET (2004): "Is Public Health Insurance an Appropriate Instrument for Redistribution?" *Discussion Paper, GREMAQ, University of Toulouse*.

KIMBALL, M. S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica*, 58, 53–73.

LOW, H. AND D. MALDOOM (2004): "Optimal Taxation, Prudence and Risk-sharing," *Journal of Public Economics*, 88, 443–464.

MIRRLEES, J. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38, 175–208.

MIYAZAKI, H. (1977): "The Rat Race and Internal Labor Markets," *Bell Journal of Economics*, 8, 394–418.

PARKER, S. C., Y. BELGHITAR, AND T. BARMBY (2005): "Wage Uncertainty and the Labor Supply of Self-employed Workers," *Economic Journal*, 115, 190–207.

PRATT, J. W. (1964): "Risk Aversion in the Small and in the Large," *Econometrica*, 32, 122–136.

RILEY, J. (1979): "Informational Equilibrium," *Econometrica*, 47, 331–359.

ROCHET, J.-C. (1991): "Incentives, Redistribution and Social Insurance," *Geneva Papers of Risk and Insurance*, 16, 143–165.

ROSS, S. A. (1981): "Some Stronger Measures of Risk Aversion in the Small and in the Large with Applications," *Econometrica*, 49, 621–638.

ROTHSCHILD, M. AND J. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay in the Economics of Incomplete Information," *Quarterly Journal of Economics*, 90, 629–649.

SANDMO, A. (1970): "The Effect of Uncertainty on Saving Decisions," *Review of Economic Studies*, 37, 353–360.

SPENCE, A. (1978): "Product Differentiation and Performance in Insurance Markets," *Journal of Public Economics*, 10, 427–447.

STIGLITZ, J. E. (1982): "Utilitarianism and Horizontal Equity," *Journal of Public Economics*, 18, 1–33.

WILSON, C. (1977): "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 12, 167–207.

Chapter 2

Bribing Thieves - A Theory of Crime and Progressive Taxation

2.1 Introduction and Literature

In the western world some provision for those threatened by the extremes of indigence or starvation due to circumstances beyond their control has long been accepted as a duty of the community [...] be it only in the interest of those who require protection against acts of desperation on the part of the needy.

F.A. von Hayek (1964, p. 285).

There is a wide-spread belief that the level of crime in highly polarized societies is larger than in more equal societies, which is not in the interest of the rich.[1] Therefore, measures have to be taken, namely redistribution from top to bottom, to prevent the 'needy' from engaging in appropriative actions. This idea is a particular aspect of the more general notion that "social peace" is a necessary requirement for the functioning of societies.[2]

However, the properties of the respective redistribution scheme are much less obvious. F.A. von Hayek seems to have had some kind of means-tested basic income in mind. Yet, means-tested basic income *lowers* incentives to pick up work and may increase the incentive to spend the time trying to appropriate others' goods or resources. Therefore, large parts of the later literature on conflict and redistribution have focussed on the question whether lump-sum transfers work as law enforcement instruments. Brennan (1973) first mentions that "appeasement policies", that increase the standard of living of poor individuals, can reduce their incentive to engage in "revolutionary activities".[3] Imrohoroglu et al. (2000) and Demougin and Schwager (2000) and (2003) formalize similar ideas in a context where property rights are costly to enforce because individuals can decide to become thieves. In their models, government raises an income tax and spends the revenue on direct law enforcement and on transfers. These transfers are withdrawn if an individual is caught committing a crime. Hence, though the different papers

[1] The word "crime" is in fact not the most accurate word to describe the conflict that is concerned in this chapter, since only theft motivated by the goal to increase consumption is considered. Still, "crime" and "theft" will be used synonymously in what follows.

[2] See, for example, the narrative arguments in Mulligan and Sala-i-Martin (1999, p. 28) or Buchanan (1975, p. 85).

[3] Acemoglu and Robinson (2000) pick up this idea and argue that the extension of the franchise in many countries during the nineteenth century might have been a device for commitment to enduring redistribution, and hence a measure to prevent revolution.

examine quite different questions[4], they share the mechanism by which redistribution affects crime levels. It goes back to Becker's (1968) fundamental insight that, holding fixed the probability of being caught red-handed, a larger fine increases the expected cost of criminal activities. Paying a transfer that convicted criminals are deprived of works exactly in this way, i.e. it increases the punishment. This approach, however, leaves open some important questions. First, imposing more drastic punishments directly should work as well and will probably be preferred by the working citizens.[5] Second, effective transfers would have to be substantially larger than the subsistence level, which has to be granted even to criminals. This is usually not observed in reality. Finally, the approach explains only the *level* of redistribution but has to take the *shape* of the whole tax schedule as given. As Acemoglu and Robinson (2000) put it: "...there are also major differences in the form of redistribution across countries. [...] It is important to understand what might cause these differences..." (p.1194).

The main contribution of the present work is to offer an explanation why the whole shape of the tax schedule matters. It will be shown that direct progression – observed very frequently in reality – can help the well-endowed citizens achieve their egoistic goals, because it reduces the cost of crime prevention. In addition, the low-skilled citizens, who might be thieves before redistribution, will also profit from progressive taxation.

Arguments for why optimal crime policies can involve a distortion of individual labor supply decisions at the margin have already been given elsewhere. Grossman (1995) finds that a wage subsidy given to workers from capital owners can be in the interest of the latter. Such efficiency wages create several benefits to the employer. First, they directly increase incentives to work and decrease incentives to engage in "class-struggle". Second, higher wages induce an income effect on the

[4]Imrohoroglu et al. (2000) examine the mix between transfers and expenditures for direct law enforcement in a political equilibrium. Demougin and Schwager (2000) are interested in the cost-minimizing mix between these instruments. Comparing their results to the optimal level of redistribution and crime when transfers are based on altruism, Demougin and Schwager (2003) argue that policies in the United States might be based on "security arguments" while being based on altruism in Europe.

[5]In a model of strategic interaction between two players, Kolmar and Bös (2003) assume that any kind of direct law enforcement and hence punishment is prohibitively costly, such that redistribution is the only remaining instrument. Benoît and Osborne (1995) argue that the risk of punishing innocent citizens might put an upper bound on the optimal size of punishments.

demand for leisure and reduce the willingness of workers to spend their free time in extralegal activities. Third, the growing overall labor supply raises the interest rate in Grossman's general equilibrium model of a closed economy. The model is based, however, on some specific assumptions that make it hard to transfer the results to questions of optimal government policies in the presence of crime.[6] Most importantly, Grossman assumes that capital owners do not have the possibility to enforce their property rights directly, by investing effort into defense against criminals. The indirect way via higher wages is their only means of law enforcement, whereas any government will be able to decide on an optimal mix between direct law enforcement and redistribution.

In a paper that is closely related to the present contribution, Grossman and Kim (2003) actually show that the possibility to enforce property rights directly has an important impact on the optimal choice of other policies. They argue that governments with different law enforcement abilities will choose different educational policies. A government that is able to enforce property rights directly will favor an egalitarian educational policy, while an elitist policy will be chosen when the instrument of public law enforcement is not available. The present contribution is based on a producer-predator model of the same structure.[7] Attention is restricted to the case where the government's law enforcement ability is not exogenously constrained, and the model is solved completely under this assumption.[8] To investigate how different types of individuals are actually affected by redistribution, the optimal amount of law enforcement *without* the possibility of redistribution is determined first. It turns out that government will only find it optimal to prevent crime completely if the ratio of high-skilled to low-skilled individuals in the society is sufficiently small, but not otherwise.[9] The main proposition then states

[6] Among these drawbacks is, for example, the fact that the interest rate effect vanishes in small open economies, and that financing the wage subsidy is assumed not to cause any distortions.

[7] This type of model goes back to Grossman (1998). It is assumed that individuals have to decide whether to become "producers" or "predators", i.e. it is not possible to earn some legal *and* some illegal income. The underlying conflict situation is modeled as a contest, as suggested by Tullock (1980).

[8] Grossman and Kim (2003) impose the restriction that the ratio of high-skilled to low-skilled individuals is small. As will turn out below, the case with only few low-skilled individuals is especially interesting for the following result.

[9] This result has also been obtained by Grossman (1998), who is not concerned with redistribution.

that the introduction of progressive taxation and an appropriate adjustment of direct law enforcement effort can yield a Pareto improvement whenever the ratio of high-skilled to low-skilled individuals is not too small. This includes cases where complete prevention of crime was optimal before redistribution, but also *all* situations in which it was optimal to let the low-skilled individuals be thieves. In the first case, progressive taxation acts as a *substitute* for direct law enforcement effort, which can be reduced in response to redistribution. In the second case, progressive taxation and direct law enforcement are *complementary*, because it is optimal to accept no crime after redistribution. The Pareto improvement result is especially interesting for this latter case, because former thieves might well suffer from an increased intensity of law enforcement. Also, the result shows that the occurrence of crime is not Pareto efficient.

The remainder of the chapter is organized as follows. Section 2.2 presents the basic model and a preliminary result. Section 2.3 derives the optimal amount of law enforcement effort and the optimal level of crime in the absence of redistributive taxation. Section 4 introduces the possibility to redistribute income and contains the central result. Section 2.5 concludes. Some proofs are relegated to the Appendix.

2.2 The Model

Assume there are two types of individuals, differing in their productivity: low-skilled workers with productivity $w_l > 0$ and high-skilled workers with productivity $w_h > w_l$. When working, each individual supplies one unit of time and is paid according to its productivity. Also assume that all individuals are risk-neutral.

The number of low-skilled workers is normalized to unity. The ratio of high-skilled to low-skilled workers is denoted by $\alpha > 0$. Finally, denote the share of low-skilled workers that decide not to work but instead become thieves with $0 \leq s_l \leq 1$ and the corresponding share of high-skilled workers with s_h. This formulation presumes that each person *either* works *or* is a thief. It is not possible to split total available time between working and stealing. The total number of criminals is then equal to $s_l + \alpha s_h$. While leisure is not chosen endogenously in the model, it will become clear that income taxation can still be distortionary.

If becoming a criminal, each person uses one unit of time trying to appropriate the income of working people. It is assumed that both low- and high-skilled individuals have the same "productivity" in stealing, which can then be normalized to one without loss of generality.[10]

Modeling contests where many agents try to steal from many other, possibly heterogeneous agents can be complicated, as it is necessary to specify "...the nature of the matching process involved in agents' challenging the initial claims of other agents." (Grossman, 2001, p. 349). To keep the analysis tractable some simplifying assumptions have to be made. First, assume that law enforcement can only be provided publicly by government, financed through a constant marginal tax $0 \leq t < 1$ on gross income of workers.[11] Therefore, the total amount of income available for law enforcement is equal to

$$R(s_l, s_h, t) = t[(1 - s_l)w_l + \alpha(1 - s_h)w_h]. \tag{2.1}$$

It is assumed that this tax revenue is not subject to appropriation by thieves, either because it is raised directly in form of law enforcement services, or because material goods used for law enforcement (such as police cars, court rooms, and surveillance cameras) are not attractive to thieves.

Second, thieves cannot distinguish workers according to their productivity and therefore spend equal amounts of their appropriative effort on each of them. Thus, the same *share* of net income is stolen from every worker.[12] Given that the total amount of appropriative effort is

$$A(s_l, s_h) = s_l + \alpha s_h \tag{2.2}$$

[10]It is straightforward to modify the model to allow for differing criminal skills a_l and a_h. Under mild assumptions, the results are unchanged. First, $w_h > w_l$ has to be replaced by the corresponding "comparative advantage" assumption $w_h/a_h > w_l/a_l$. Second, the assumption that $w_l > 1$ below is rewritten as $w_l > a_l$. Finally, the assumption that all thieves capture the same amount of goods must be replaced by the assumption that the loot is distributed in proportion to criminal ability.

[11]This assumption implies that property can neither be defended privately nor individually-targeted by the government.

[12]Another interpretation of the contest environment would be to say that thieves are randomly matched to workers. In this case, one has to interpret utilities as expected utilities, which does not make a difference in the present setup but would considerably complicate the analysis with risk-averse agents.

and using a standard Tullock contest success function, this share is equal to

$$p = \frac{A(s_l, s_h)}{A(s_l, s_h) + R(s_l, s_h, t)} \tag{2.3}$$

whenever $A(s_l, s_h) > 0$, and $p = 0$ otherwise.[13] The consumption of a working individual with productivity w_i is therefore equal to

$$c_i^w = w_i(1-t)(1-p). \tag{2.4}$$

Per capita consumption of thieves is derived by dividing the total amount of stolen income by the number of thieves. It is given by

$$c^s = \frac{p(1-t)[(1-s_l)w_l + \alpha(1-s_h)w_h]}{A(s_l, s_h)} \tag{2.5}$$

whenever $A(s_l, s_h) > 0$, i.e. whenever thieves exist. The time structure of the model is as follows. First, government sets a tax rate. Then the equilibrium shares s_l and s_h are determined. Workers then produce income and thieves spend their appropriative effort. Produced income is distributed according to (2.3).

The model is solved backwards. Given a tax rate t, the shares s_l and s_h can be endogenized. Since everyone has the opportunity to switch from being a criminal to being a worker, the incomes c_i^w and c^s have to be equal if the equilibrium share is $0 < s_i < 1$. If $c_i^w > c^s$ ($c_i^w < c^s$) for all small (large) values of s_i, then $s_i = 0$ ($s_i = 1$) will be an equilibrium.

Lemma 2.1. *Given a tax rate $0 \le t < 1$, the equilibrium share of thieves among individuals with productivity $w_i,\ i = l, h$, is given by*

$$s_i^* \in \begin{cases} \{0\} & \textit{if} \quad t > 1/w_i \\ [0,1] & \textit{if} \quad t = 1/w_i \\ \{1\} & \textit{if} \quad t < 1/w_i. \end{cases}$$

Proof. See Appendix, Section 2.A.1. □

Some interesting implications follow immediately. First, if the tax rate is smaller than $1/w_h$ nobody will work. This is intuitively reasonable: even if all high-skilled individuals worked, such small tax rates would not generate enough revenue for

[13]While appropriation and guarding efforts enter (2.3) without relative weights, the relative efficiency of law enforcement can still be varied by changing the parameters w_l and w_h of the model, holding the ratio w_h/w_l fixed.

law enforcement to make it profitable to be a worker. Second, if the low-skilled individuals have a very low productivity $w_l \leq 1$, they can never be induced to become workers. This rather unappealing case will be excluded in the following by the assumption that $w_l > 1$. Finally, if $t = 1/w_i$ holds for some i, any share of thieves within group i can occur in equilibrium, while the equilibrium share within in the other group will be uniquely determined. However, as will become clear when the government's objective is discussed in the next section, a situation with $t = 1/w_i$ and an interior value for s_i will never be optimal. Government will always have an incentive to raise the tax rate slightly above $1/w_i$ to get rid of criminals at negligible additional costs for the previous workers. To circumvent technical complications associated with such considerations, it will be assumed from now on that $s_i^* = 0$ if $t = 1/w_i$. The effect of the tax rate t can then be summarized as follows. For low tax rates $t < 1/w_h$ nobody produces, with a resulting consumption of zero for everyone (case I). As the tax rate is increased, high-skilled workers start to produce beginning with $t = 1/w_h$. For rates $1/w_h \leq t < 1/w_l$ the low-skilled workers are still thieves (case II), but higher tax rates reduce the stolen share p. If the tax reaches $t = 1/w_l < 1$, low-skilled workers also start to produce (case III).

2.3 The Optimal Level of Crime

2.3.1 *The Government's Problem*

Deriving an optimal tax rate and therefore an optimal level of crime requires the specification of government's objective. For example, government might want to maximize per-capita consumption of the high-skilled individuals. This approach is used in comparable papers such as Demougin and Schwager (2000) and (2003) and Grossman and Kim (2003). It is especially appealing because the later case for redistribution is strongest if based on an objective that incorporates no considerations of equity whatsoever.

Alternatively, the maximization of *workers'* per-capita consumption could be the goal of government. This objective can be justified from a political economy perspective, assuming that criminals do not participate as political citizens. A problem, however, arises from the fact that the group of worker-voters is endogenous and reacts to the tax rate as outlined in the previous section. Depending on the exact specification of the political process, more than one political equilibrium could exist. In the present context, maximization of workers' per-capita consumption is not a well-defined objective.

Finally, government might be interested in maximizing the overall level of consumption or production. Moreover, the objective of production maximization is equivalent to an aversion to crime as a matter of principle. In both cases government will want to circumvent any theft irrespective of what share of production must be invested in law enforcement.

Throughout this section it will be assumed that the tax rate t is chosen as to maximize per-capita consumption of the high-skilled. The robustness of the results with respect to the other objectives is discussed in section 2.3.5.

Observe that the government's objective is not continuous, since raising the tax rate to $1/w_i$ induces a drop of s_i from one to zero. This problem will be solved as follows. First, the optimal tax rate will be determined for each of the three cases outlined above. The resulting optimal consumption levels of the high-skilled are then compared to find the global optimum.

2.3.2 *Case I: No Production*

From the previous section it is clear that a tax rate $0 \leq t < 1/w_h$ is not optimal for *any* of the discussed government objectives. Raising it to $t = 1/w_h$ makes everybody better off since production starts at this level.

2.3.3 *Case II: Production and Theft*

Set $s_l = 1$ and $s_h = 0$. Using these conditions together with (2.3) and (2.4) yields consumption of the high-skilled as:

$$c_h^w(t) = \frac{t(1-t)\alpha w_h^2}{1+\alpha t w_h}. \tag{2.6}$$

The first derivative of (2.6) with respect to t is given by

$$\frac{\partial c_h^w}{\partial t} = \frac{\alpha w_h^2[1-2t-\alpha t^2 w_h]}{(1+\alpha t w_h)^2}. \tag{2.7}$$

Inspection of the second derivative reveals that c_h^w is strictly concave in t throughout. The first-order condition of the problem therefore characterizes an optimal tax rate $\tilde{t}^*$. Furthermore, if $\tilde{t}^* \leq 1/w_h$, the optimal rate of the case II problem, t^*, is given by the corner solution $t^* = 1/w_h$. For values $1/w_h < \tilde{t}^* < 1/w_l$ the optimal tax rate t^* is equal to $\tilde{t}^*$. If $\tilde{t}^* \geq 1/w_l$, a solution to the problem does not exist, but the value $c_h^w(1/w_l)$ corresponds to the supremum of (2.6) under case II. From the first order condition (2.7),

$$\tilde{t}^* = \frac{1}{1+\sqrt{1+\alpha w_h}}. \tag{2.8}$$

Setting $\tilde{t}^* \leq 1/w_h$ and simplifying yields

$$\alpha \geq w_h - 2 =: \overline{\alpha}. \tag{2.9}$$

Hence, if the ratio of high-skilled to low-skilled workers is sufficiently large, it will be optimal for a government that is restricted to case II to impose the lower-bound tax rate $t^* = 1/w_h$.[14] In such societies, relatively few thieves cause too little damage to spend any additional effort on law enforcement. Equivalently, solving $\tilde{t}^* \geq 1/w_l$ yields

$$\alpha \leq \frac{w_l}{w_h}\left[w_l - 2\right] =: \underline{\alpha}. \tag{2.10}$$

Observe that $\underline{\alpha} < \overline{\alpha}$ whenever $\overline{\alpha}$ is actually positive. If the ratio of high-skilled to low-skilled workers is sufficiently small, government will want to impose a tax rate as high as possible, because in such societies the harm that criminal low-skilled do to the high-skilled workers is large and justifies high levels of law enforcement.

To summarize, a government that maximizes per-capita consumption of the high-skilled and is restricted to case II chooses the low tax rate $t^* = 1/w_h$ if $\alpha \geq \overline{\alpha}$ and a rate as close as possible to $1/w_l$ if $\alpha \leq \underline{\alpha}$. For $\underline{\alpha} < \alpha < \overline{\alpha}$ it chooses $t^* = \tilde{t}^*$ as given in (2.8). It can easily be shown that this rate is strictly decreasing and convex in α.

Before examining the next case it is necessary to derive the high-skilled consumption supremum under case II, denoted by c^{II}. This can later be compared to consumption under complete deterrence of crime. After substitution of the optimal tax rate as described above and some simplifications

$$c^{II} = \begin{cases} \frac{w_h^2 \alpha (w_l - 1)}{w_l(\alpha w_h + w_l)} & \text{if} \quad \alpha \leq \underline{\alpha} \\ \frac{2 + \alpha w_h - 2\sqrt{\alpha w_h + 1}}{\alpha} & \text{if} \quad \underline{\alpha} < \alpha < \overline{\alpha} \\ \frac{\alpha(w_h - 1)}{1+\alpha} & \text{if} \quad \alpha \geq \overline{\alpha}. \end{cases} \tag{2.11}$$

[14]If $w_h \leq 2$ this will always be the case, because $\overline{\alpha} \leq 0$ while $\alpha > 0$.

Some properties of (2.11) deserve further attention. First, c^{II} is strictly increasing in α for all values of $\alpha > 0$, given that $1 < w_l < w_h$. It is also continuous in α. It is sketched in Figure 1 for the case that $\underline{\alpha} > 0$, together with some results from the next section. The three functions that define c^{II} stepwise as given in (2.11) are represented as dotted curves. The function c^{II} itself is highlighted by a solid line.

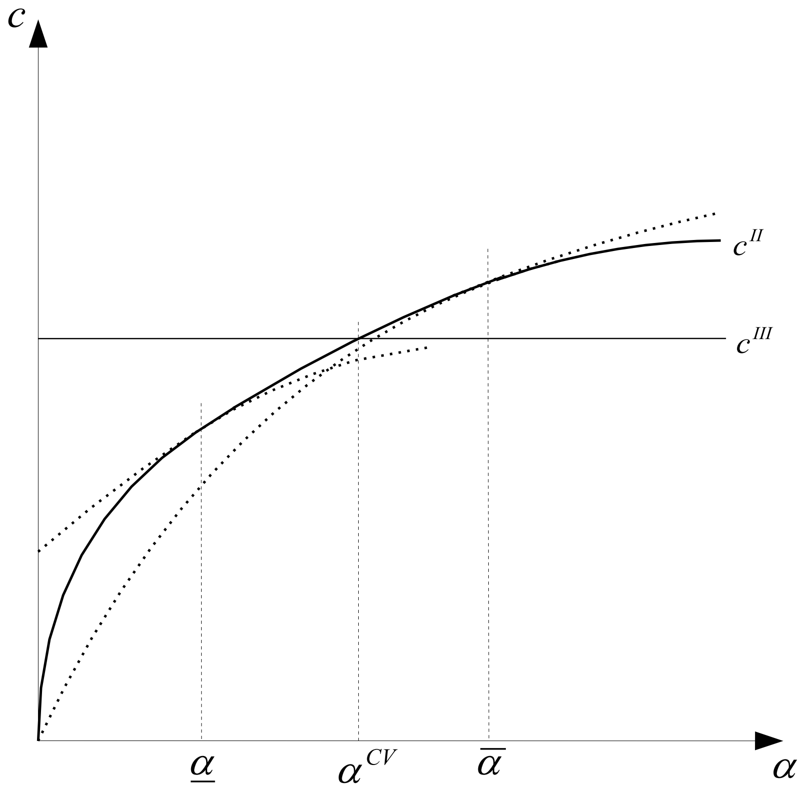

Figure 2.1: Consumption under optimal taxation

2.3.4 *Case III: No Theft*

Setting the tax rate to $1/w_l$ deters even the low-skilled individuals from engaging in criminal activities. Larger tax rates can never be optimal, since there is no more crime to prevent and any additional law enforcement effort would be a waste of resources. Consumption of a high-skilled individual is then given by

$$c^{III} = w_h \left(1 - \frac{1}{w_l}\right), \tag{2.12}$$

due to $p = 0$. This is evidently independent of α, and hence c^{III} is depicted as a horizontal line in Figure 1. Together with the results from the previous section a critical level α^{CV} follows, from which on the high-skilled workers' consumption is larger under optimization within case II than under case III, and vice versa. It is obtained by setting expression (2.12) equal to (2.11). To begin with, use c^{II} for α in between $\underline{\alpha}$ and $\overline{\alpha}$. The resulting critical value is

$$\alpha^{CV} = 4\frac{w_l}{w_h}(w_l - 1). \tag{2.13}$$

Given that $w_l > 1$, expression (2.13) is positive and also larger than $\underline{\alpha}$.[15] It is, however, possible that (2.13) is also larger than $\overline{\alpha}$, in which case (2.13) does not provide the correct critical value. The correct one is then derived by setting (2.12) equal to c^{II} for α larger than $\overline{\alpha}$. This gives

$$\alpha^{CV} = \frac{w_h}{w_h - w_l}(w_l - 1). \tag{2.14}$$

2.3.5 *Summary and Robustness*

Proposition 1 summarizes the results from the previous subsections.

Proposition 2.1. *There exists a unique critical value $\alpha^{CV} > \underline{\alpha}$, given by either (2.13) or (2.14), so that government that maximizes per-capita consumption of the high-skilled chooses the law enforcement tax rate*

$$t^{opt} = \begin{cases} 1/w_l & \text{if} \quad \alpha < \alpha^{CV} \\ \tilde{t}^* & \text{if} \quad \alpha \geq \alpha^{CV} \wedge \alpha < \overline{\alpha} \\ 1/w_h & \text{if} \quad \alpha \geq \alpha^{CV} \wedge \alpha \geq \overline{\alpha}, \end{cases} \tag{2.15}$$

where $\tilde{t}^*$ is given by (2.8).

It will only be optimal to completely prevent crime if there are few high-skilled individuals, since only the tax rate $1/w_l$ is large enough to deter all potential criminals. This is optimal as otherwise the large number of criminal low-skilled individuals would do much harm to the working high-skilled. Overall, the optimal tax rate (weakly) decreases in the ratio of high- to low-skilled individuals.

[15] The fact that $\alpha^{CV} > \underline{\alpha}$ shows that the range of (2.11) where it does not give the maximum consumption but the supremum is never globally optimal.

How do the results change if government uses a different objective function? Of course, for maximization of production or with complete crime aversion the optimal rate would always be equal to $1/w_l$, so that all potential criminals are deterred and become workers instead. The picture is less clear-cut if total consumption, i.e. aggregate production less law enforcement effort was to be maximized, independent of its distribution among workers and thieves. Starting from a situation in which Proposition 1 prescribes an interior tax rate $t^{opt} = \tilde{t}^*$, decreasing the tax rate down to $1/w_h$ reduces law enforcement effort, affecting the distribution of production but not its aggregate level. Hence total consumption goes up. On the other hand, setting t up to $1/w_l$ induces an upwards jump in production, which might overcompensate for the larger levels of law enforcement effort. The optimal tax rate would therefore be equal to either $1/w_l$ or $1/w_h$, depending on the specific parameter constellations. Qualitatively, results remain similar. With many low- compared to high-skilled individuals, the gain of many extra workers will be decisive, making the large rate $1/w_l$ optimal.

2.4 Crime and Redistribution

So far, taxes are only raised to finance law enforcement. This section is concerned with the question whether government can do even better by additionally redistributing income.

Most models discussed in the introduction assume that government can distinguish honest and criminal individuals with positive probability, with law enforcement affecting this probability. Individuals can then be treated conditional on their observed status. Uncovered criminals receive punishment, while all others receive a positive transfer. The difference between these two treatments works as the effective punishment, and larger transfers as well as stronger direct punishments reduce crime levels.

The mechanism in Grossman and Kim (2003) differs from this approach. The complex process of crime and law enforcement is completely summarized in the success function (2.3). In this framework it would be inconsistent to assume that policies can directly condition on an individual's status of being criminal. It then becomes clear that neither means-tested minimal income nor lump-sum transfers to all individuals can help in accomplishing the government's goal. Minimal income is awarded to citizens who cannot claim enough legal income, the thieves among them. It therefore lowers incentives to work legally and is counterproductive. A lump-sum transfer to all individuals, financed through a larger tax rate, also reduces returns to work and does not make crime less attractive. In this sense the model highlights an effect which often seems to be ignored in the public discussion about inequality and crime. It is certainly not a trivial truth that any measure aimed at equalizing standards of living will reduce incentives to engage in appropriative actions and hence crime levels.[16]

A different redistribution scheme is more promising, namely a positive marginal tax on high incomes and a negative marginal tax on low incomes, which amounts to the introduction of progressive taxation. This measure makes work more attractive to low-skilled individuals, and less law enforcement effort is required to prevent them from being thieves. Hence progressive taxation lowers the cost of complete crime deterrence. If complete deterrence was optimal even without such redistribution, a mix of direct law enforcement effort and redistribution may be less costly to the high-skilled than law enforcement alone. In this case, progressive taxation serves as a *substitute* for direct law-enforcement, because law enforcement effort can be reduced in response to redistribution. If complete deterrence was not optimal without redistribution, it may become optimal because redistribution reduces its cost. Redistribution is then a *complement* to law enforcement, which becomes stricter as taxation becomes progressive.

In any case, the search for the optimal redistribution scheme will take place under the assumption that nobody steals. With the low-skilled remaining criminals, there is no sensible effect of progressive taxation because nobody will claim the "wage subsidy". The utility levels with optimal redistribution and complete prevention of crime will then have to be compared to the utility levels without redistribution but optimally chosen law enforcement expenditures as given in the previous section.

[16]There exist, however, model modifications under which minimal income could again have positive effects. These modifications could include non-linearities in the utility functions, such as a subsistence level of income. Guaranteeing this minimal standard might alleviate the problem of theft. Satiation might play a similar role.

Assume that government does not accept any theft, i.e. it sets the tax rate to $1/w_l$, which depends inversely on the low-skilled wage. If this wage can be increased, law enforcement effort can be saved. This establishes the trade-off that serves as justification for progressive taxation. Consider an additional tax $\tau > 0$ which is imposed on high incomes. Assuming gross income to be the tax base of τ, the "new productivity" of the high-skilled workers becomes

$$\hat{w}_h(\tau) = w_h(1-\tau). \tag{2.16}$$

As before, government can collect this tax before individual income is subject to possible appropriation effort by thieves. Put differently, the assumption that only individuals' *net* income can be stolen is maintained, possibly because taxes are deducted at source. A difference to taxes generated for law enforcement arises, however, when it comes to expenditure. Assume that the revenue generated by τ is paid to the low-skilled by means of a subsidy $\gamma(\tau)$ on gross income, so that

$$\hat{w}_l(\tau) = w_l(1+\gamma(\tau)) \tag{2.17}$$

is their new productivity after redistribution. Redistribution raises the low-skilled workers' income, and the received transfers are thus potentially exposed to theft. The relation between τ and γ is established by the government's budget constraint

$$\gamma(\tau) = \frac{\alpha w_h}{w_l}\tau. \tag{2.18}$$

Clearly, the additional restriction that $\hat{w}_l(\tau) \leq \hat{w}_h(\tau)$ has to be imposed, i.e. redistribution should not reverse the order of incomes, since high-skilled individuals would otherwise decide to become thieves. This defines a maximal redistributive rate $\overline{\tau} = (w_h - w_l)/(w_h(1+\alpha)) > 0$, at which $\hat{w}_l(\overline{\tau}) = \hat{w}_h(\overline{\tau}) = (\alpha w_h + w_l)/(1+\alpha)$ holds.

For a given rate τ, the new rate $\hat{t}(\tau)$, whose revenue is still exclusively spent on law enforcement, is now equal to $1/\hat{w}_l(\tau)$, which decreases in τ. Starting from the initially linear tax schedule, the introduction of τ and $\gamma(\tau)$ amounts to the introduction of progression. Workers' net incomes become

$$\hat{c}_h^w(\tau) = w_h(1-\tau)(1-\hat{t}(\tau)) \text{ and } \hat{c}_l^w(\tau) = w_l(1+\gamma(\tau))(1-\hat{t}(\tau)). \tag{2.19}$$

Hence the effective tax rate on high incomes is $\hat{t}+\tau-\hat{t}\tau$ but only $\hat{t}+\hat{t}\gamma-\gamma$ on low incomes.[17] As the following proposition establishes, the introduction of progression can lead to a Pareto improvement in a large range of cases. In deriving the result it is assumed that the government chooses the redistributive tax rate that again maximizes utility of the high-skilled. This assumption is not restrictive here, because it turns out that the low-skilled are never made worse off with this approach.[18]

Proposition 2.2. *If $\alpha > (w_l/w_h)(w_l - 1)$, the introduction of the optimal redistributive rate for the high-skilled,*

$$\tau^* = \min\left\{\frac{\sqrt{\alpha w_h + w_l} - w_l}{\alpha w_h}, \overline{\tau}\right\} > 0, \tag{2.20}$$

accompanied by an adjustment of t^{opt} to $1/\hat{w}_l(\tau^*)$, yields a Pareto improvement. If $\alpha \leq (w_l/w_h)(w_l - 1)$, progressive taxation cannot be Pareto improving.

Proof. See Appendix, Section 2.A.2. □

Observe first that the critical value $(w_l/w_h)(w_l - 1)$ is smaller than α^{CV}, irrespective of whether α^{CV} is given by (2.13) or (2.14). Hence progressive taxation accompanied by complete prevention of crime can always be Pareto improving when it was optimal to accept some crime before redistribution. This fact is interesting for two reasons. First, the Pareto improvement result is especially surprising in this case, where redistribution and law enforcement act as complements. The word "complements" so far has been used to indicate that the introduction of progressive taxation is accompanied by complete prevention of crime. As can easily be shown, the overall *effort* invested in law enforcement does also increase.[19] It is therefore not obvious that former thieves will also profit from the "carrot and stick" policy change. Second, it implies that the acceptance of crime as derived in section 2.3 is

[17] After substitution of (2.18), the effective rate on low incomes turns out to be smaller than the rate on high incomes whenever $\hat{t} < 1$, which is satisfied. Note that the effective rate on low incomes might both be positive or negative.

[18] If the *optimal* rate for the high-skilled did not lead to a Pareto improvement because the low-skilled were made worse off, a Pareto improvement might still be possible by choosing a different, possibly larger rate.

[19] This is trivially fulfilled if $\alpha \geq \overline{\alpha}$, so that the low law enforcement tax rate $1/w_h$ was optimal before redistribution. In that case, the law enforcement rate is raised to $1/\hat{w}_l > 1/w_h$ and the number of tax payers increases as well. The statement also holds if $\alpha^{CV} \leq \alpha < \overline{\alpha}$, so that the law enforcement rate might decrease from $\tilde{t}^*$ to $1/\hat{w}_l$ but the number of tax payers still increases in equilibrium.

never Pareto efficient once government has the instrument of progressive taxation available in addition to the option of direct law enforcement.

If α and/or w_h are small, progression will not be desirable from point of view of the high-skilled. This holds due to the influence of α and w_h on the budget constraint (2.18). If there are too many low-skilled workers, who in turn would have to be subsidized by relatively few high-skilled, the instrument of redistribution is ineffective in lowering the cost of crime prevention.

Consider finally the optimal rate τ^*. As shown in the proof of Proposition 2.2, there exists a unique value $\widetilde{\alpha} > \overline{\alpha}$ so that the maximal rate $\overline{\tau}$ will be optimal if and only if $\alpha \geq \widetilde{\alpha}$. Hence if the share of low-skilled individuals is small enough, it will be optimal to redistribute until all incomes in the society are equalized. This is another interesting implication of the above proposition. It again holds due to the budget constraint: if α is large, even small rates τ already facilitate large subsidies γ, making redistribution an effective instrument to lower the cost of direct law enforcement.

2.5 Conclusions

This contribution analyzes a model of theft in heterogeneous societies. It first solves the problem of optimal law enforcement when redistribution of income through progressive taxation is not an available policy instrument. It is shown that the acceptance of some crime is optimal when the share of high-skilled individuals in the society is sufficiently large. Otherwise, any crime should be deterred. It then turns out that the introduction of progressive taxation, accompanied by an appropriate adjustment of direct law enforcement effort, can yield a Pareto improvement under a large variety of circumstances, because redistribution lowers the costs of keeping the low-skilled from engaging in appropriative actions. The model therefore provides an explanation not only for some level of redistribution in the presence of social conflict, but is able to justify a directly progressive shape of the whole tax schedule based on pure efficiency arguments.

In particular, a Pareto improvement is always possible if it was optimal to accept crime before redistribution. In this case, progression and direct law enforcement act as complements, because no crime should be accepted after redistribution. Both workers and former thieves will profit from this policy reform. The acceptance of crime is therefore not efficient.

A weakness of the model is that the redistributive tax induces only a discrete distortion. Excessive redistribution would lead high income earners to stop working and switch to criminal activities. The welfare improving effect of redistribution might therefore be overstated.[20] Still, the fundamental idea remains valid in a model that allows for a labor-leisure trade-off or a continuous division of time between production and appropriation. The same holds for a model extension that would allow for more than two types; the intuition of the results does not hinge on this assumption.

A useful extension would consider different matching of criminals to workers as well as the possibility to privately invest in the defense of property rights. Working might be more profitable for the low-skilled if they are targeted by less intense appropriation effort than the rich. Within an intertemporal extension of the model with uncertainty, the effects of public policy instruments such as social security and social insurance may finally generate interesting insights.

[20] One should keep in mind, however, that labor supply is highly more elastic on the extensive margin than on the intensive margin (see e.g. Heckman (1993)). This evidence suggests that the present model may well be appropriate, despite its restriction on labor supply decisions.

2.A Appendix

2.A.1 Proof of Lemma 2.1

The result follows directly from deriving the difference $\Delta c_i = c_i^w - c^s$ using (2.1)-(2.5). It is defined whenever $0 < s_i < 1$ and given by

$$\Delta c_i = \frac{(1-t)(tw_i - 1)[\alpha(1-s_h)w_h + (1-s_l)w_l]}{s_l + \alpha s_h + t[\alpha(1-s_h)w_h + (1-s_l)w_l]}. \tag{2.21}$$

First, the denominator of (2.21) is positive whenever $0 < s_i < 1$. The sign of (2.21) is therefore equal to the sign of its numerator. If $t = 1$, then $\Delta c_i = 0$ since any possible production is spent for law enforcement. Any profile of shares would then be an equilibrium. This irrelevant case has been excluded by the assumption that $t < 1$. The sign of the numerator is then equal to the sign of $tw_i - 1$. Specifically, this holds irrespective of the share of criminals in the own and in the other group.

2.A.2 Proof of Proposition 2.2

Given τ, set $\hat{t}(\tau) = 1/\hat{w}_l(\tau)$ to completely deter crime. Consumption of a high-skilled worker is then

$$\hat{c}_h^w(\tau) = \hat{w}_h(\tau)\left(1 - \frac{1}{\hat{w}_l(\tau)}\right). \tag{2.22}$$

Using (2.16), (2.17) and (2.18), it is easy to show that $\hat{c}_h^w(\tau)$ is strictly concave in τ. Maximizing (2.22) with respect to τ, without any additional restrictions, yields an optimal rate

$$\widetilde{\tau}^* = \frac{\sqrt{\alpha w_h + w_l} - w_l}{\alpha w_h}. \tag{2.23}$$

Clearly, $\widetilde{\tau}^* > 0$ if and only if $\alpha > (w_l/w_h)(w_l - 1)$. Therefore, if $\alpha \leq (w_l/w_h)(w_l - 1)$ progressive taxation is not part of an optimal policy mix and a Pareto improvement is impossible. Observe that the critical value $(w_l/w_h)(w_l - 1)$ is smaller than α^{CV}, irrespective of whether α^{CV} is given by (2.13) or by (2.14).

Assume that $\alpha > (w_l/w_h)(w_l - 1)$. Concerning the other possible corner solution, the condition that $\widetilde{\tau}^* \geq \overline{\tau}$ can be rearranged to

$$\sqrt{\alpha w_h + w_l} \leq 1 + \alpha. \tag{2.24}$$

This is clearly not satisfied at $\alpha = (w_l/w_h)(w_l - 1)$, where $\widetilde{\tau}^* = 0$ and $\overline{\tau} > 0$. On the other hand, it will be satisfied for α sufficiently large, since $\lim_{\alpha\to\infty} \sqrt{\alpha w_h + w_l}/(1+\alpha) = 0$. Given that the RHS of (2.24) is increasing linearly in α while the LHS is continuous and strictly concave in α, it follows that there must be a unique value $\widetilde{\alpha} > (w_l/w_h)(w_l - 1)$ so that $\widetilde{\tau}^* \geq \overline{\tau}$ if and only if $\alpha \geq \widetilde{\alpha}$. The critical value $\widetilde{\alpha}$ is the largest value that solves (2.24) as an equality, and is given by $\widetilde{\alpha} = \frac{1}{2}(w_h - 2) + \frac{1}{2}\sqrt{(w_h - 2)^2 + 4(w_l - 1)}$. Observe that $\widetilde{\alpha} > \overline{\alpha} = w_h - 2$ because $w_l > 1$.

Assume first that $\alpha < \widetilde{\alpha}$, so that $\widetilde{\tau}^*$ from (2.23) is the optimal redistributive tax rate from point of view of the high-skilled, under complete deterrence of crime. Then,

$$\hat{c}_h^w(\widetilde{\tau}^*) = \frac{1 + \alpha w_h + w_l - 2\sqrt{\alpha w_h + w_l}}{\alpha}. \tag{2.25}$$

The consumption of the low-skilled workers is given by

$$\hat{c}_l^w(\widetilde{\tau}^*) = \sqrt{\alpha w_h + w_l} - 1. \tag{2.26}$$

Consider now the case in which complete deterrence of crime is optimal even without redistribution, i.e. the case $\alpha < \alpha^{CV}$. Clearly, high-skilled workers profit from progression whenever $\widetilde{\tau}^* > 0$. In this case, low-skilled workers are better off as well. They profit from both the subsidy $\gamma(\widetilde{\tau}^*) > 0$ and the lower crime enforcement rate $\hat{t}(\widetilde{\tau}^*) = 1/\hat{w}_l(\widetilde{\tau}^*)$.

Consider next the case $\alpha^{CV} \leq \alpha < \overline{\alpha}$, where complete deterrence of crime is not optimal without redistribution. The value of (2.25) now has to be compared to the respective consumption before redistribution from (2.11),

$$c^{II} = \frac{2 + \alpha w_h - 2\sqrt{\alpha w_h + 1}}{\alpha}. \tag{2.27}$$

The condition $\hat{c}_h^w(\widetilde{\tau}^*) > c^{II}$ can be rewritten as

$$w_l - 1 > 2\left(\sqrt{\alpha w_h + w_l} - \sqrt{\alpha w_h + 1}\right). \tag{2.28}$$

Set $\alpha = 0$ and rearrange to obtain $(w_l - 1)^2 > 0$, which is satisfied because $w_l > 1$. Next observe that the RHS of (2.28) is decreasing in α, so that (2.28) is satisfied for all values of $\alpha > 0$. As for the low-skilled, their consumption under the optimal regime from section 3 is given by

$$c^s = \sqrt{\alpha w_h + 1} - 1, \tag{2.29}$$

which is strictly smaller than (2.26) because $w_l > 1$.

Finally, consider the case $\overline{\alpha} \leq \alpha \wedge \alpha^{CV} \leq \alpha$. We now need to compare

$$c^{II} = \frac{\alpha(w_h - 1)}{1 + \alpha} \tag{2.30}$$

from (2.11) to (2.25). Examine the difference $\Delta c_h = \hat{c}_h^w(\widetilde{\tau}^*) - c^{II}$, which can be rearranged to

$$\Delta c_h = \frac{w_h + \alpha}{1 + \alpha} + \frac{1 + w_l - 2\sqrt{\alpha w_h + w_l}}{\alpha}. \tag{2.31}$$

Note first that (2.31) is strictly increasing in w_l because $\alpha w_h + w_l > 1$. Hence, to show that the whole expression (2.31) will always be positive, set w_l as small as possible, i.e. examine the case where $w_l = 1$. Solving $\Delta c_h = 0$ for α then yields $\alpha = \overline{\alpha}$ as the unique solution. Since $\alpha \geq \overline{\alpha}$, $w_l > 1$, Δc_h is continuous in α (where $\alpha > 0$) and $\lim_{\alpha \to \infty} \Delta c_h = 1$, it follows that (2.31) is strictly positive within the given range of α. Consumption of the low-skilled thieves before redistribution is given by

$$c^s = \frac{\alpha(w_h - 1)}{1 + \alpha}. \tag{2.32}$$

The corresponding difference $\Delta c_l = \hat{c}_l^w(\widetilde{\tau}^*) - c^s$ can be rearranged to obtain

$$\Delta c_l = \sqrt{\alpha w_h + w_l} - \frac{\alpha w_h + 1}{1 + \alpha}. \tag{2.33}$$

As before, this expression is increasing in w_l, so the case where $w_l = 1$ is examined. Solving $\Delta c_l = 0$ for α again gives $\overline{\alpha}$ as a solution, in addition to 0 and $-1/w_h$. Continuity and the fact that $\lim_{\alpha \to \infty} \Delta c_l = +\infty$ imply that $\Delta c_l > 0$ for the relevant parameter constellations.

Assume now that $\alpha \geq \widetilde{\alpha}$, so that $\overline{\tau}$ is the optimal redistributive tax rate for the high-skilled under complete deterrence of crime, and incomes are equalized. Consumption of both types of individuals is then given by

$$\hat{c}(\overline{\tau}) = \frac{\alpha(w_h - 1) + (w_l - 1)}{1 + \alpha}. \tag{2.34}$$

Again consider the case where $\alpha < \alpha^{CV}$. The fact that the introduction of $\overline{\tau}$ together with a reduction of t from $1/w_l$ to $1/\hat{w}_l(\overline{\tau})$ is Pareto improving follows as above when $\alpha < \widetilde{\alpha}$ was assumed to hold.

The case $\alpha^{CV} \leq \alpha < \overline{\alpha}$ cannot occur together with $\alpha \geq \widetilde{\alpha}$ because $\widetilde{\alpha} > \overline{\alpha}$.

Finally, consider again the case $\overline{\alpha} \leq \alpha \wedge \alpha^{CV} \leq \alpha$. Consumption of both types under the optimal regime without redistribution is given by the identical expressions (2.30) or (2.32), which are obviously strictly smaller than (2.34).

References of Chapter 2

ACEMOGLU, D. AND J. ROBINSON (2000): "Why Did the West Extend the Franchise? Democracy, Inequality and Growth in Historical Perspective," *Quarterly Journal of Economics*, 115, 1167–1199.

BECKER, G. (1968): "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 76, 169–217.

BENOÎT, J.-P. AND M. OSBORNE (1995): "Crime, Punishment and Social Expenditure," *Journal of Institutional and Theoretical Economics*, 151, 326–347.

BRENNAN, G. (1973): "Pareto Desirable Redistribution: The Non-Altruistic Dimension," *Public Choice*, 14, 43–67.

BUCHANAN, J. (1975): "The Political Economy of Franchise in the Welfare State," in *Capitalism and Freedom: Problems and Prospects*, ed. by R. Selden, Charlottesville: University Press of Virginia, 52–91.

DEMOUGIN, D. AND R. SCHWAGER (2000): "Excess Burden of Criminality and Redistribution," *International Review of Law and Economics*, 20, 329–351.

——— (2003): "Law Enforcement and Criminality: Europe vs. USA," *International Review of Law and Economics*, 23, 217–225.

GROSSMAN, H. (1995): "Robin Hood and the Redistribution of Income," *European Journal of Political Economy*, 11, 399–410.

——— (1998): "Producers and Predators," *Pacific Economic Review*, 3, 169–187.

——— (2001): "The Creation of Effective Property Rights," *American Economic Review, Papers and Proceedings*, 91, 347–352.

GROSSMAN, H. AND M. KIM (2003): "Educational Policy: Egalitarian or Elitist?" *Economics & Politics*, 15, 225–246.

HECKMAN, J. (1993): "What Has Been Learned About Labor Supply in the Past Twenty Years?" *American Economic Review*, 83, 116–121.

IMROHOROGLU, A., A. MERLO, AND P. RUPERT (2000): "On the Political Economy of Income Redistribution and Crime," *International Economic Review*, 41, 1–25.

KOLMAR, M. AND D. BÖS (2003): "Aggressiveness and Redistribution," *Bonn Econ Discussion Papers*, 2/2003.

MULLIGAN, C. AND X. SALA-I-MARTIN (1999): "Social Security in Theory and Practice (II): Efficiency Theories, Narrative Theories, and Implications for Reform," *NBER Working Paper Series*, 7119.

TULLOCK, G. (1980): "Efficient Rent Seeking," in *Toward a Theory of the Rent-Seeking Society*, ed. by J. Buchanan, R. Tollison, and G. Tullock, College Station: Texas A&M University Press, 97–112.

VON HAYEK, F. (1964): *The Constitution of Liberty*, Chicago: University of Chicago Press.

CHAPTER 3

Evolution of Time Preferences and Attitudes Towards Risk

3.1 Introduction and Literature

This chapter is based on the hypothesis that individual decisions are guided by hedonic utility. An individual who faces several alternatives will choose the one that promises the greatest pleasure, or happiness. Then, given that the properties of a hedonic utility function determine individual behavior, and individual behavior determines biological fitness, evolutionary forces will have shaped our utility during the long time in which the modern human being evolved.

In this sense, utility can be considered as a “reward system” that induces individuals to make evolutionary optimal choices, a view which is supported both by theory and evidence from neuroscience (Damasio (1994), Bechara et al. (1997), Gazzaniga et al. (1998)). As Kupfermann et al. (2000, p. 1007) express it: “Pleasure is unquestionably a key factor in controlling the motivated behaviors of humans”. For the economist, the interesting question is then about the properties of the evolutionary optimal reward system, and how these properties adapt to the environment in which individuals make choices.

The present work reconsiders and solves a general model of the evolution of utility suggested by Robson (2001a), which predicts how cardinal properties of utility functions should adapt to the decision environment. It turns out that the optimal utility function will be steep in regions where decisions have to be made frequently, and where wrong decisions would lead to large losses in fitness. In those regions, even small changes in consumption will cause large changes in happiness.[1]

The general model suggests several applications. In the context of intertemporal preferences, it predicts that individuals should be especially impatient concerning waiting times which they face often. Under the assumption that more immediate judgements are necessary more frequently, the immediate future should be discounted at a higher rate than the distant future. This explains why additional waiting time causes large discomfort in the short-run but much less in the long-run. As a corollary, one should expect patience to increase as children grow up and have to make more far-reaching decisions.

[1] In the present contribution, evolution of utility is examined in a non-strategic environment. Koçkesen et al. (2000a) and (2000b) and Dekel et al. (2007) model the evolution of preferences in strategic environments. Koçkesen et al. (2000b) and (2000b) identify classes of games in which negatively interdependent preferences can lead to an advantage for the respective players. Dekel et al. (2007) explore the role of observability of preferences for the evolution of subjective payoffs in games.

The simultaneous evolution of several different utility functions is then shown to be optimal whenever the decisions which our ancestors had to make arrived in distinct choice situations. In the context of time preferences, this may have led to the development of different utility functions for short-run and for long-run decisions. The model therefore provides an evolutionary justification for the "multiple-selves" approach to time discounting (see Frederick et al. (2002) for an overview), where dynamic inconsistency arises from a conflict between different decision mechanisms. This view has been corroborated by the results of McClure et al. (2004), who show that different parts of the brain are active in short-run and in long-run decisions. The present model predicts that conflict between the "myopic" and the "farsighted" mechanism is more likely to occur if the decision-maker is used to small payoffs in the short-run. It also sheds light on the evolutionary role of self-commitment as introduced by Strotz (1955).

In the context of attitudes towards risk, the model highlights an influence of environmental randomness on individual preferences which has not yet been discussed by the literature. Risk attitudes will not only be influenced by the technology that converts consumption into fitness, as in Robson (1996) and (2001b), but also by the distribution according to which opportunities arise for the decision maker.[2] The model then offers an immediate evolutionary rationale for S-shaped value functions as in prospect theory (Kahneman and Tversky (1979)). Most interestingly, it identifies the individual's reference point with the peak of the density that describes the availability of alternatives. This provides a clear prediction of the reference point even in highly stochastic environments.

The contributions by Rayo and Becker (2007a) and (2007b) also deal with the evolution and adaption of hedonic utility. In their model, optimal happiness derived from income is a step function with a unique jump, which can be interpreted as the aspiration level an agent wants to achieve. The aspiration level can then be shown to adjust over time and in accordance with income levels of a peer group, given that payoffs are correlated over time and across individuals. This offers an evolutionary explanation for habit formation and peer comparisons, phenomena

[2]The literature, including Cooper and Kaplan (1982), Robson (1996), Bergstrom (1997) and Curry (2001) has also highlighted the role of aggregate risk, which makes deviations from standard expected utility maximization evolutionarily optimal.

frequently observed in happiness surveys.[3] The present contribution addresses different questions, making use of a different model. While the underlying adaption mechanisms share similarities – in the sense that utility adjusts to the decision environment like an eye to the ambient brightness (Frederick and Loewenstein (1999)) – the model outlined below derives utility as a tool to make reasonable comparisons between any pair of alternatives, as opposed to identifying only the best out of a large set. For the analysis of time preferences and attitudes towards risk, this turns out to be an appropriate starting point.

The chapter is organized as follows. The general model and its solution are presented in section 3.2. Section 3.3 is devoted to the evolution of intertemporal preferences. Section 3.4 is concerned with attitudes towards risk. Section 3.5 concludes. A longer proof can be found in the Appendix.

3.2 A General Model

3.2.1 Description

The model in this section has been suggested by Robson (2001a). It has been solved there for an approximate evolutionary criterion, the *probability of mistakes* criterion. In the following, it will be solved under the correct objective, the *expected loss* criterion.

Assume an agent repeatedly has to make choices between alternatives from a set $X = [a, b]$, which are identified with fitness. Thus, alternative $x \in X$ yields fitness x, where fitness can simply be thought of as the number of offspring.[4] When making a decision, the agent does not face the whole set X, but only two alternatives that are independently drawn from X according to the same random distribution. The agent has to choose one of these alternatives. Assume that the random distribution can be represented by a bounded density f with finitely many discontinuities. Denote the corresponding distribution function by F.

The distribution represents the agent's environment by describing the availability of different alternatives. For example, during good times, in fertile geographical regions, or under a favorable climate, large fitness alternatives will available

[3] Samuelson (2004) shows that relative consumption effects can be an evolutionarily optimal way for the decision-maker to utilize information about the state of nature contained in the consumption of others.

[4] The case where alternatives are not directly equated with fitness levels is considered later in this section.

with greater probability than otherwise. Changes in the environment can later be modeled through changing distributions. For the moment, the distribution is considered as fixed.

The agent is endowed with a hedonic utility function that assigns a level of pleasure or happiness to each element in X. The alternative that promises larger pleasure will be chosen. The question now is: which utility function leads to the largest expected fitness? It is motivated by the idea that evolution will eventually have "discovered" and selected this optimal function.

Without any restrictions on the set of admissible functions, the problem is trivial. Any strictly increasing utility function ensures that the better of any two alternatives will correctly be identified. This is, on the other hand, not a realistic assumption. Happiness cannot be perceived in arbitrarily fine shades, due to limitations of human sensory abilities.[5] This constraint can be modeled by assuming that utility can only take discrete, albeit extremely many values. In the following, the set of admissible utility functions is thus restricted to the set of increasing step functions with $N \in \mathbb{N}$ jumps of size $1/N$ each, where $0 < N < \infty$. As a result, the agent cannot distinguish two alternatives located on the same step of the utility function. Any choice between such alternatives will have to be random and a mistake can occur. Clearly, different step utility functions will then lead to different levels of expected fitness.

The size of N measures the degree of the perceptual constraint, which vanishes as $N \to \infty$.[6] The assumption that each jump corresponds to a utility increment of size $1/N$ ensures that utility is normalized to the interval $[0, 1]$ for all values of N. In the following, results will be derived for the limiting case where $N \to \infty$, motivated by the presumption that perceptual constraints do exist but are extremely small. Also, the optimal limiting utility function turns out to be continuous. It is thus an easy-to-deal-with approximation for a step function with a huge number of steps.

The problem of finding the optimal step utility function is equivalent to the problem of locating N thresholds in the set X, where two alternatives can only be distinguished if there is at least one threshold between them.[7] Whenever two al-

[5] See, for example, Kupfermann et al. (2000, p. 421 ff.).

[6] The assumption of a large but finite N is analogous to the limited perception constraint by Rayo and Becker (2007a) and (2007b).

[7] This representation of the problem is in fact the one used by Robson (2001a).

ternatives are drawn from in between two neighboring thresholds, the agent will choose the worse one with probability $1/2$. Robson (2001a) has analyzed the problem of locating the thresholds to minimize the probability of such mistakes, obtaining a simple and intuitive solution, which will be replicated below. The appropriate evolutionary criterion, however, is the maximization of expected fitness, or, equivalently, the minimization of the expected loss due to wrong decisions.

3.2.2 Solution

The problem is solved here in three steps. First, the density f is approximated by a sequence of step densities with finitely many steps, in a way that ensures uniform convergence to f as the number of steps grows to infinity. Second, the problem of optimal threshold location is considered for each of these step densities, yielding utility functions for the limit as $N \to \infty$.[8] Finally, the behavior of these functions is examined as the step densities converge to f.

Assume without loss of generality that f is left-continuous and consider a step density $\hat{f}_S$ that approximates f as follows. Let $y_i, i = 1, ..., D-1$, be the points where f is discontinuous, and define $y_0 := a$ and $y_D := b$, where $D \geq 1$. Hence X can be partitioned into D intervals on which f is continuous. Each of these intervals is then decomposed into $S \geq 1$ steps of equal length, so that there are $S \cdot D$ steps altogether. For $i = 1, ..., SD-1$, let $\pi(i) = \lfloor i/S \rfloor$ be the largest integer smaller or equal to i/S. Then define $x_0 := a$, $x_{SD} := b$ and for each $i = 1, ..., SD-1$,

$$x_i := y_{\pi(i)} + (i - \pi(i) \cdot S) \left(\frac{y_{\pi(i)+1} - y_{\pi(i)}}{S} \right).$$

Now let $X_0 := \emptyset$ and for $i = 1, ..., SD$ define $X_i := \{x \in X | x \leq x_i\} \setminus \bigcup_{j=0}^{i-1} X_j$. Clearly, f is continuous on each step $X_i, i = 1, ..., SD$. Denote by $L(X_i) := x_i - x_{i-1}$ the length of step X_i. Now define

$$\hat{f}_S := \sum_{i=1}^{SD} \mathbb{I}_{X_i} f_i$$

[8] It is easy to show that an optimal solution to the problem exists for any number of thresholds N. Let $T = \{t \in [a, b]^N | a \leq t_1 \leq ... \leq t_N \leq b\}$ be the domain of the optimization, where t_k denotes the position of the k-th smallest threshold in $[a, b]$. Clearly, T is compact. Since the loss function as defined in the Appendix is continuous, the statement follows from the Weierstrass Theorem. Should there be several solutions, the following results hold for any selection of them.

where $\mathbb{I}_{X_i}$ is the indicator function of X_i, and $f_i := (1/L(X_i)) \int_{x_{i-1}}^{x_i} f(y)dy$ is a value taken by $f(x)$ somewhere on X_i (by continuity of f on X_i), which makes sure that $\hat{f}_S$ is again a density. It also follows that $\hat{f}_S$ converges uniformly to f on X as $S \to \infty$.[9]

To obtain the utility of an alternative $x \in X$ for a fixed profile of N thresholds, the number of thresholds below x has to be multiplied by $1/N$. Denote then by $\theta_{N,S}(x)$ the number of thresholds below x given that N thresholds have been located to maximize expected fitness under the step density $\hat{f}_S$. The resulting utility is given by $U_{N,S}(x) = \theta_{N,S}(x)/N$. For comparison, let $\vartheta_{N,S}(x)$ be the number of thresholds below x if the probability of mistakes is minimized, yielding utility $V_{N,S}(x) = \vartheta_{N,S}(x)/N$. The main result of this section can now be stated as follows:

Theorem 3.1. *For each $x \in X$*

$$V(x) := \lim_{S\to\infty} \lim_{N\to\infty} V_{N,S}(x) = \int_a^x f(y)dy = F(x),$$

and

$$U(x) := \lim_{S\to\infty} \lim_{N\to\infty} U_{N,S}(x) = c \int_a^x f(y)^{2/3}dy,$$

where $c = (\int_a^b f(y)^{2/3}dy)^{-1}$ is a normalizing constant.

Proof. See Appendix, Section 3.A.1. □

The limiting utility function $V(x)$, which follows from minimizing the probability of mistakes, equals the distribution function $F(x)$. The same result has been obtained by Robson (2001a), who solves for the optimal threshold positions directly.[10] Intuitively, when only the mistake probability is concerned, many thresholds should be allocated to regions of X where decisions have to be made with large probability, i.e. where the density $f(x)$ is large. Avoiding mistakes in this region is particularly beneficial. The limiting utility function will then be steep in this region, resembling the distribution function.

[9] This is a direct implication of the fact that f is continuous and bounded on each of the D intervals defined above.

[10] When the mistakes probability is minimized, the problem can be solved without the detour via step densities. This approach is not transferable to the case of loss minimization.

Evolution maximizes expected fitness, for which the size of mistakes matters as well. As shown in the Appendix, the expected size of a mistake between two thresholds depends on the cube of the distance between them. For this reason, strong variations in the distances between thresholds will be undesirable, making the evolutionary optimal distribution of thresholds more even than indicated by the first result. In particular, this implies that the slope of $U(x)$ will not vary as much as the slope of $F(x)$, which is achieved by the concave transformation of $f(x)$ in the definition of $U(x)$.[11]

Besides Theorem 3.1, the threshold model delivers an intuitive interpretation for the slope of a utility function. Since a large slope derives from a dense allocation of thresholds, one can think of the local steepness of a utility function as the degree of attention devoted to this area. The curvature properties of utility then correspond to changes in attention.

3.2.3 Extensions

The analysis above proceeded under the assumption that choices are made between fitness levels directly. Assume instead that the individual makes choices between alternatives from a set $Y \subseteq \mathbb{R}$, which are not identified with fitness levels. In a decision situation, two alternatives are independently drawn from Y according to a distribution function G. They are then mapped to fitness by a function $\psi : Y \rightarrow [a, b]$, which again induces a distribution of fitness levels in X, for which an optimal utility function can be derived. The utility assigned to alternative $y \in Y$ becomes $U(\psi(y))$, which will simply be denoted by $U(y)$ (or $V(y)$, respectively) with some abuse of notation.

For Theorem 3.1 to be applicable, the induced distribution of fitness levels needs to to be representable by a bounded density $f(x)$ with finitely many discontinuities. This requirement is not very restrictive and can be ensured by various different joint assumptions on Y, ψ and G. For example, make the following (strong) assumption.

Assumption 3.1. *$Y = [d, e]$, ψ is continuously differentiable with $\psi'(y) > 0$ for all $y \in Y$, $\psi(d) = a$ and $\psi(e) = b$, and the distribution on Y can be represented by a continuous density g.*

[11] Clearly, $U(x)$ and $V(x)$ coincide for a uniform fitness distribution.

Assumption 3.1 is by no means necessary for the theorem to be applicable, but it ensures that the induced fitness density $f(x)$ is continuous, producing continuously differentiable utility functions. The utility functions defined on Y are

$$U(y) = c \int_a^{\psi(y)} f(x)^{2/3} dx$$

for the expected loss criterion and

$$V(y) = F(\psi(y)) = G(y)$$

for the probability of mistakes criterion. Closer inspection of $U(y)$ reveals the following result:

Proposition 3.1. *Under Assumption 3.1, the function $U(y)$ is continuously differentiable with*

$$U'(y) = c\, g(y)^{2/3} \psi'(y)^{1/3}.$$

$U(y)$ is twice differentiable at each $y \in Y$ where $g(y)$ and $\psi'(y)$ are differentiable and $g(y) > 0$, with

$$U''(y) = \frac{2}{3} c \left(\frac{\psi'(y)}{g(y)} \right)^{1/3} g'(y) + \frac{1}{3} c \left(\frac{g(y)}{\psi'(y)} \right)^{2/3} \psi''(y).$$

Proof. Derive the induced fitness distribution function $F(x)$ first. Since ψ is strictly increasing, $F(x) = G(\psi^{-1}(x))$ clearly holds. $F(x)$ is continuously differentiable under Assumption 3.1, with derivative $F'(x) = f(x) = g(\psi^{-1}(x))(\partial \psi^{-1}(x)/\partial x) = g(\psi^{-1}(x))/\psi'(\psi^{-1}(x))$. Hence

$$U'(y) = c\, \psi'(y) f(\psi(y))^{2/3} = c\, g(y)^{2/3} \psi'(y)^{1/3},$$

from which the statement about $U''(y)$ follows immediately. □

The proposition shows that the slope of $U(y)$ at y corresponds to a normalized weighted geometric mean of the slope of the distribution function $G(y)$ and of the slope of the fitness function $\psi(y)$. The utility function $U(y)$ therefore represents an intermediate case between the actual fitness function $\psi(y)$ (properly normalized) and the distribution function $G(y)$. Intuitively, utility should again be steep in regions of Y where decisions have to be made often. However, since the size of mistakes – measured in fitness – matters as well, it should also inherit properties of the fitness function ψ. Specifically, when ψ is steep somewhere, thresholds should be spaced closely there, because a wrong decision is severely damaging even if the two alternatives at choice are very close to each other. On the other hand, mistakes are not very damaging in regions where ψ is almost flat and all alternatives yield very similar fitness levels.

The second derivative of $U(y)$, if it exists, corresponds to a weighted average of $g'(y)$ and of $\psi''(y)$.[12] If, for example, the fitness function $\psi(y)$ is concave, utility can still be convex if $G(y)$ is convex.

So far, the assumption of a one-dimensional set of alternatives $Y \subseteq \mathbb{R}$ has been made. The model can be extended to higher-dimensional sets $Y \subseteq \mathbb{R}^n$, though, as long as there is a single-valued fitness function $\psi : Y \to [a, b]$ and a distribution function G on Y that induces a bounded fitness density $f(x)$ with finitely many discontinuities.[13] As above, the optimal utility from an alternative $y = (y_1, ..., y_n) \in Y$ then becomes

$$U(y_1, ..., y_n) = c \int_a^{\psi(y_1,...,y_n)} f(x)^{2/3} dx.$$

Under the respective differentiability conditions, the following partial derivatives can be obtained:

$$\frac{\partial U(y)}{\partial y_i} = c \frac{\partial \psi(y)}{\partial y_i} f(\psi(y))^{2/3}$$

for all $i = 1, ..., n$. It then immediately follows that for all $i, j = 1, ..., n$, the marginal rates of substitution between y_i and y_j according to $U(y)$ will coincide with those according to $\psi(y)$; indifference curves will coincide with fitness isoquants. The partial derivatives of $U(y)$ and of $\psi(y)$, however, while having the same sign, will not in general coincide. Second partial derivatives can also be examined. Omitting the dependency on y for notational simplicity, it holds for all $i = 1, ..., n$ that

$$U_{ii} = cf(\psi)^{2/3}\psi_{ii} + \frac{2}{3} c \psi_i^2 f(\psi)^{-1/3} f'(\psi),$$

[12] Observe that the weights change with y. Whenever $g(y)$ is small compared to $\psi'(y)$, the curvature of $G(y)$ is relatively more important. When $\psi'(y)$ is large, the curvature of $\psi(y)$ matters more.

[13] The combination of a multi-dimensional set of alternatives and a single-valued fitness function is referred to as a "fitness landscape" in biology (Wright (1932)).

where indices denote partial derivatives. Hence the local curvature properties of $U(y)$ are again influenced by those of the fitness function and of the distribution function.

3.2.4 Hedonic Adaption

If our ancestors' environment as described by G varied in a systematic way during a human lifetime, evolution should have selected individuals whose utility functions accommodate to change.[14] An *adaption mechanism* can be thought of as a set of different utility functions together with a rule that specifies which function will become active at what point in time. In general, adaption of utility will have to be triggered by perceivable changes in the environment, which were (and might still be) correlated with changes in the density G, which is itself not directly observable. For example, an accumulation of large payoffs will generally indicate that the environment has developed in a favorable way.[15] Realized payoffs are, however, not the only possible trigger for hedonic adaption. If, for example, the nature of decision problems changed systematically with individual age or with the seasons, utility should be expected to differ between age groups or between summer and winter. Some of these different triggers will be discussed in the following.

3.3 Intertemporal Preferences

3.3.1 A Simple Model of Impatience

To apply the results of section 3.2 to intertemporal preferences, start with a very simple model with continuous time, where the agent has to choose between prospects that yield a fitness payoff normalized to 1 after a waiting time of $t \geq 0$, which starts right after the decision has been made. Waiting times are drawn from $Y = [0, T]$ according to a continuous density g.

Assume that the fitness of an alternative with waiting time t, evaluated at the point in time where the decision is made, is given by the exponential function $\psi(t) = e^{-\delta t}$, where $\delta > 0$ is a constant discount factor. There are various reasons for discounting delayed payoffs in such a way. If, for example, there is a constant hazard that the payoff vanishes while the agent waits for it, as in Sozou

[14] The case of changing fitness functions ψ or changing sets of alternatives Y will be taken up in section 3.3.

[15] This is essentially the adaption mechanism at work in Rayo and Becker (2007a) and (2007b), where realized payoffs contain information about environmental change due to correlation of payoffs across time and individuals.

(1998) or Dasgupta and Maskin (2005), the expected fitness of an alternative with waiting time t is calculated as above, where δ is the hazard rate of the underlying failure distribution. Alternatively, population growth is a reason for discounting, as shown by Hansson and Stuart (1999) and Robson and Samuelson (2007).

Setting $X = [a, b] = [e^{-\delta T}, 1]$ ensures that a modified version of Assumption 3.1 applies, where $\psi' > 0$ is replaced by $\psi' < 0$, and $\psi(0) = b$ and $\psi(T) = a$ holds. It then immediately follows that $U(t)$ is continuously differentiable with

$$U'(t) = -c\, g(t)^{2/3}(-\psi'(t))^{1/3} < 0.$$

The derivative of $U(t)$ is a local indicator for impatience, since it measures the effect on happiness of having to wait a little longer. The two most interesting questions are then how $U'(t)$ is affected by the decision environment in general, and specifically how impatience varies with waiting time t. A natural reference point for such considerations could be the impatience of a person that discounts exponentially. The condition that $U'(t) < \psi'(t)$, i.e. that $U(t)$ exhibits greater impatience than $\psi(t)$ at t, can easily be rearranged to $c^{3/2} g(t) > -\psi'(t)$. Consider, for example, a case where $g(t)$ is very large for small values of t but decreases quickly. It then follows that impatience will be larger for small waiting times but smaller for long waiting times, compared to impatience from exponential discounting. The described situation is one where choices between short-run alternatives are more frequent than between long-run alternatives, which is likely to be an adequate description of our ancestors' environment. Compared to today, for example, hunter-gatherers' planning horizons were undoubtedly shorter, due to lack of information and prediction possibilities, technologies for storage, and possibly shorter life expectancy. The general insights of section 3.2 hence apply to impatience in a straightforward way: if more immediate judgements are necessary more frequently, the near future will be discounted at a higher rate than the remote future.

Now consider a model where the optimal utility function $U(t)$ can be derived explicitly. Assume that $Y = [0, \infty)$, $X = [0, 1]$ and that waiting times are drawn according to an exponential distribution with $g(t) = \lambda e^{-\lambda t}$ and $\lambda \geq \delta > 0$.[16] The

[16]The assumption that $\lambda \geq \delta$ is necessary for the induced fitness density to be bounded in the present model where waiting times are not bounded.

parameter λ can be used to vary the decision environment. Larger values of λ imply that more probability mass is concentrated on short waiting times.

Proposition 3.2. *If waiting times are drawn from* $[0,\infty)$ *according to the density* $g(t) = \lambda e^{-\lambda t}$ *with* $\lambda \geq \delta$, *then* $U(t) = e^{-\gamma t}$, *where* $\gamma = \frac{2}{3}\lambda + \frac{1}{3}\delta$.

Proof. Proceeding as in the proof of Proposition 3.1 yields the induced fitness density $f(x) = \frac{\lambda}{\delta}x^{(\frac{\lambda}{\delta}-1)}$, which is bounded whenever $\lambda \geq \delta$. It then follows that c, as defined in Theorem 3.1, is given by $c = (\frac{\delta}{\lambda})^{2/3}(\frac{2\lambda+\delta}{3\delta})$. The function $U(t) = c\int_0^{\psi(t)} f(x)^{2/3}dx$ can now easily be derived explicitly, yielding the desired result. □

The proposition shows that hedonic utility will take an exponential form whenever waiting times are distributed exponentially. If $\lambda = \delta$, it follows that $\gamma = \delta$ and thus $U(t) = \psi(t) = e^{-\delta t}$. This result is based on the fact that induced fitness levels are uniformly distributed when $\lambda = \delta$. If $\lambda > \delta$, i.e. small waiting times and hence larger fitness levels occur relatively more often, the future is discounted at the larger rate $\gamma > \delta$.[17]

As a final remark, note that the results presented in this initial section cannot yet explain behavior such as preference reversals and dynamic inconsistency. After all, optimal utility will be increasing in the fitness of an alternative.[18] The results can, however, be tested by assessing self-reported measures of impatience and how these measures change over time. It is reasonable to assume, for example, that hunter-gatherer children did generally not have to make many long-run decisions, but increasingly more of them as they grew up. In the context of the above approach, this could be modeled as a decrease of λ during adolescence. The reader can now form an own opinion about the prediction that impatience should decline over the life-cycle, at least up to a certain age.

3.3.2 The Discrete Time Model

The following simple model allows alternatives to differ both in waiting time t and in fitness payoff v, which realizes after the waiting time has passed.

[17] Consider again the local impatience measure $U'(t) = -\gamma e^{-\gamma t}$. For $\lambda > \delta$ it follows that $U'(t) < \psi'(t)$ if $t < t^*$ and $U'(t) > \psi'(t)$ if $t > t^*$, where $t^* = (\ln\gamma - \ln\delta)/(\gamma - \delta) > 0$. This again illustrates that short-run impatience will be larger than under standard discounting if short-run decisions have to be taken relatively more often.

[18] An exception would be the case where the density $g(t)$ is zero on a range of waiting times, yielding a flat utility function there. Preference reversals could then occur based on indifference in the initial decision.

Assume that time is discrete and both waiting times and payoffs are bounded, such that $Y = \{0, 1, ..., T\} \times [0, 1]$. The fitness of an alternative $y = (t, v) \in Y$ is given by $\psi(t, v) = \delta^t v$, where $0 < \delta < 1$ is again a discount factor.[19] Alternatives are drawn from Y as follows. First, a waiting time $t \in \{0, 1, ..., T\}$ is drawn according to strictly positive probabilities $p_0, p_1, ..., p_T$. Conditional on having drawn t, a payoff $v \in [0, 1]$ is chosen according to a distribution function $G_t(v)$ with continuous density $g_t(v)$. Issues related to returns on investment can be captured by the assumption that the densities g_t change with t in a systematic way. This setup will be referred to as the "discrete time model" in the following.

Lemma 3.1. *In the discrete time model, fitness levels are distributed in $X = [0, 1]$ according to the density*

$$f(x) = \sum_{t=0}^{\hat{t}(x)} \frac{p_t}{\delta^t} g_t \left(\frac{x}{\delta^t} \right),$$

where $\hat{t}(x)$ is the largest waiting time t for which $x \leq \delta^t$ holds.

Proof. Conditional on having drawn $t \in \{0, 1, ..., T\}$, the fitness levels are distributed in $[0, \delta^t]$ according to the distribution function $F_t(x) = G_t(x/\delta^t)$. Unconditionally, fitness levels are thus distributed in $[0, 1]$ according to the distribution function

$$F(x) = \sum_{t=0}^{\hat{t}(x)} p_t \, G_t \left(\frac{x}{\delta^t} \right) + \sum_{t=\hat{t}(x)+1}^{T} p_t,$$

where $\hat{t}(x)$ is the largest waiting time t for which $x \leq \delta^t$ still holds. For waiting times larger than $\hat{t}(x)$, even the best attainable fitness level will be smaller than x. It then follows immediately that $F(x)$ is differentiable everywhere except (possibly) at the points $x = \delta^t$ for $t = 1, ..., T$, with derivative $f(x)$ as given in the Lemma. □

Observe that the fitness density $f(x)$ is left-continuous with possible downwards jumps at the points $x = \delta^t$ for $t = 1, ..., T$. A jump occurs because slightly larger fitness levels than x are no longer attainable with a waiting time of $\hat{t}(x)$ if $x = \delta^t$, i.e. $\hat{t}(x)$ is a step function with downwards jumps at $x = \delta^t$, $t = 1, ..., T$. Still, continuity of $f(x)$ at $x = \delta^t$ will hold if $\lim_{v \to 1} g_t(v) = 0$.

[19] As in the previous section, the discount factor represents the probability that the reward vanishes in any given period, or that the decision-maker dies.

A modified version of the discrete time model will be used in the next section, where it provides a simple framework to analyze dynamic inconsistency. Section 3.4 will also revert to it.

3.3.3 Dynamic Inconsistency

A choice situation as described above consists of a set of alternatives Y, a fitness function ψ, and a distribution function G on Y. From this, a utility function $U : Y \to [0, 1]$ can be obtained. Hedonic adaption has been interpreted as an adjustment of utility to a changing environment as described by G. Apparently, choice situations can also differ in the set of available alternatives Y. Hunter-gatherers are frequently confronted with typical hunt decisions, involving the choice between different hunting strategies (which animal to hunt, which technique to use). The choice between different foraging strategies appears as a different decision problem, which can clearly be distinguished from the first one. Yet another choice situation will have involved the long-run choice between different areas of habitation.

If individuals are capable of distinguishing these situations, evolution should have endowed them with different utility functions, each one tailored to a specific problem, and activated by the recognition of the respective choice situation. The same assumption has been made in the context of hedonic adaption, where utility changes over time. In the present context it is more convenient to envisage the simultaneous existence of several choice mechanisms, which generate optimal decisions in their respective area of specialization.

Consider the following modified version of the discrete time model. Assume that an agent is faced with two possible choice situations. The first one involves short-run alternatives with waiting time $t = 0$ and payoffs v that are drawn from $[0, 1]$ according to the density g_0. The second situation involves alternatives with waiting times $t \in \{1, 2\}$, where $t = 1$ occurs with probability $0 < p < 1$ and $t = 2$ occurs with probability $1 - p$. Payoffs are drawn from $[0, 1]$ conditional on waiting time according to the densities g_1 and g_2. The type of decision situation (short-run or long-run) is revealed to the agent before the actual choice, so that the optimal utility function can be activated. The crucial difference to the standard model from the previous section is therefore that the function $U(0, v)$ evolves independently from the function $U(t, v)$ for $t = 1, 2$. The resulting two functions, or decision mechanisms, can be interpreted as the "multiple selves" proposed by Winston (1980), Schelling (1984) and Ainslie and Haslam (1992).[20] In a recent study, McClure et al. (2004) find that there are actually two different neural systems involved in intertemporal decision-making. Functional magnetic resonance imaging reveals that the limbic system is especially active when immediate payoffs are evaluated, while the lateral prefrontal cortex is relatively more engaged in long-run decisions.

The present model thus offers an evolutionary rationale for the evolution of multiple selves, and explains "why either type of agent emerges when it does" (Frederick et al., 2002, p. 376). As will become clear in the next section, it also sheds light on the question why "farsighted selves often attempt to control the behaviors of myopic selves, but never the reverse" (ibid.). The crucial assumption behind the result is that the individual *either* faces two short-run *or* two long-run alternatives. While being a strong assumption, it is by no means less natural then the assumption that decisions have to be made between any two arbitrary alternatives, projects, or bundles of goods. It involves, however, the implicit assumption that choices are irreversible. An initial choice between two alternatives $(1, v_1)$ and $(2, v_2)$ appears as a choice between $(0, v_1)$ and $(1, v_2)$ one period later. If the initial choice could be reconsidered, short- and long-run decision would no longer be separate. Irreversibility will be assumed for the rest of this subsection.

The fitness of short-run alternatives is distributed in $[0, 1]$ according to the density $f_S(x) = g_0(x)$. In long-run decisions, fitness levels are distributed in $[0, \delta]$ according to

$$f_L(x) = \begin{cases} (p/\delta)\, g_1(x/\delta) & if \quad x > \delta^2 \\ (p/\delta)\, g_1(x/\delta) + ((1-p)/\delta^2)\, g_2(x/\delta^2) & if \quad x \le \delta^2. \end{cases}$$

The utility function used to evaluate short-run alternatives is thus given by

$$U(0, v) = c_S \int_0^v f_S(x)^{2/3} dx,$$

while

$$U(t, v) = c_L \int_0^{\delta^t v} f_L(x)^{2/3} dx$$

[20] Elster (1985) and Read (2001) model the "strategic interaction" between multiple selves explicitly.

obtains for $t = 1, 2$. Now consider two long-run alternatives:

Definition 3.1. *The decision between $y_1 = (1, v_1)$ and $y_2 = (2, v_2)$ creates regret if (i) $U(1, v_1) \leq U(2, v_2)$, and (ii) $U(0, v_1) > U(1, v_2)$.*

Regret refers to a case in which the agent initially prefers the alternative with the larger waiting time, but, after one period has passed and the earlier alternative has moved to the present, would prefer to reverse the decision.[21] The main proposition in this section states that regret will arise whenever the agent is accustomed to sufficiently small payoffs in the short run. To be able to formalize "sufficiently small", it makes use of the following definition.

Definition 3.2. *Let $h(v; \lambda)$ be a family of continuous densities on $[0, 1]$, parameterized by $\lambda > 0$, that satisfies that for all $y > 0$*

$$\lim_{\lambda \to \infty} \frac{\int_0^y h(v; \lambda)^{2/3} dv}{\int_0^1 h(v; \lambda)^{2/3} dv} = 1.$$

According to the definition, raising λ shifts probability mass to the left, in the sense that the whole relative area under the function $h(v; \lambda)^{2/3}$ concentrates below y as $\lambda \to \infty$, for any strictly positive y. This property is satisfied by several well-known distributions. Consider, for example, the truncated exponential distribution where $h(v; \lambda) = (\lambda e^{-\lambda v})/(1 - e^{-\lambda})$ and thus $(\int_0^y h(v; \lambda)^{2/3} dv)/(\int_0^1 h(v; \lambda)^{2/3} dv) = (1 - e^{-\frac{2}{3}\lambda y})/(1 - e^{-\frac{2}{3}\lambda})$. Another simple example that is in line with Definition 3.2 is the triangular distribution with $h(v; \lambda) = 2(\lambda + 1) - 2(\lambda + 1)^2 v$ if $v \leq 1/(\lambda + 1)$, and $h(v; \lambda) = 0$ otherwise.

Proposition 3.3. *Assume that $g_1(v) > 0$ for all $v \in [0, 1]$. Then for any $v_1, v_2 \in (0, 1)$ with $v_1 \leq \delta v_2$, there exists a value $\bar{\lambda}(v_1, v_2) \in \mathbb{R}$ such that the decision between $y_1 = (1, v_1)$ and $y_2 = (2, v_2)$ creates regret if $g_0(v) = h(v; \lambda)$ for any $\lambda > \bar{\lambda}(v_1, v_2)$.*

Proof. The condition that $v_1 \leq \delta v_2$, or $\psi(1, v_1) \leq \psi(2, v_2)$ immediately implies that $U(1, v_1) \leq U(2, v_2)$, which is condition (i) in the definition of regret. The assumption that $g_1(v) > 0$ for all $v \in [0, 1]$ implies that $f_L(x) > 0$ for all $x \in [0, \delta]$. This in turn implies that $0 < U(1, v_2) < 1$, because $0 < v_2 < 1$. Optimal short-run utility is given by

$$U(0, y) = \frac{\int_0^y h(v; \lambda)^{2/3} dv}{\int_0^1 h(v; \lambda)^{2/3} dv}.$$

[21] The difference between $U(0, v_1)$ and $U(1, v_2)$ could be interpreted as disutility from temptation as in Gul and Pesendorfer (2001).

According to Definition 3.2, $\lim_{\lambda\to\infty} U(0,y) = 1$ for all $y > 0$. Therefore, there exists a $\bar{\lambda}(v_1, v_2)$ such that $U(0, v_1) > U(1, v_2)$ for all $\lambda > \bar{\lambda}(v_1, v_2)$, which is condition (ii) in the definition of regret. □

If the agent is mostly confronted with small payoffs in short-run decisions, it will experience large levels of pleasure $U(0, v)$ even if v is small. This is a direct implication of the general insights derived in Section 3.2. Conflict between the farsighted and the myopic self will then occur. While the individual preferred the alternative with longer waiting time in the original decision, the earlier alternative becomes exceptionally tempting as soon as it has moved to the presence and is evaluated according to short-run utility. Note, however, that the original decision will always be optimal in an evolutionary sense.

3.3.4 Conflict and Self-Awareness

The regret considered so far is of purely seductive nature. The agent would like to reverse the initial decision, but is not able to do so by assumption. Irreversibility appears as a reasonable assumption for many day-to-day decisions in hunter-gatherer societies. In particular, many of the examples that economists refer to in the context of dynamic inconsistency involve the existence of easily transferable money, such as the premature spending of savings that were intended to finance Christmas gifts (Strotz (1955)). On the other hand, as pointed out by Laibson (1999, p. 444), "all illiquid assets provide a form of commitment", and storage is necessarily illiquid in hunter-gatherer societies. In particular, a large share of storage in primitive societies takes the form of somatic capital (Robson and Kaplan (2007)), which makes any savings decision irreversible.

Observed preference reversals are therefore likely to be the result of a maladaption to a world in which decisions have become increasingly reversible. Dasgupta and Maskin (2005) argue that reversals will be the consequence of adaption to a world in which the relative fitness of two alternatives actually changed as time went by.[22] The advantage of the present approach is that it sheds light on the phenomenon of "self-awareness" (Strotz (1955) and Pollack (1968)) and the farsighted

[22] They assume that payoffs can always realize earlier than expected. If early realization does not occur, the later alternative becomes relatively less fit. Sozou (1998) also discusses the evolution of non-exponential discounting, but the model does not explain dynamic inconsistency as considered here.

self's urge to constrain future choices of the myopic self. Such commitment makes long-run decisions irreversible and thus allows the maintenance of specialized decision mechanisms. One should therefore expect the emergence of self-awareness as a substitute in a world of increasing reversibility. From there, the coexistence of multiple selves and the ability to foresee future choice inconsistencies is not as paradoxical as it may appear: different utility functions make better choices possible, and self-constraints prevent wrong utility functions from interfering later.

As already mentioned at the end of the last section, the initial decision between alternatives $(1, v_1)$ and $(2, v_2)$ is always optimal regarding fitness. It is then clear that the farsighted self should try to constrain the myopic self, but evolution should not endow the myopic self with a desire to constrain the farsighted self.

3.4 Attitudes Towards Risk

Now turn to the evolution of attitudes towards risk. Some first implications for the curvature properties of evolutionarily optimal utility functions follow immediately. Assume that individuals choose between fitness levels from $X = [a, b]$ directly, which are drawn and offered according to the density $f(x)$. It follows immediately that both $U(x)$ and $V(x)$ will be concave, i.e. exhibit risk-aversion globally, if $f(x)$ is strictly decreasing. This corresponds to a situation in which choices between alternatives with small payoffs have to be made more often than between alternatives with large payoffs.

It also follows that $U(x)$ will in general exhibit less risk-aversion than $V(x)$, i.e. its slope varies less than that of $V(x)$. This can most easily be seen by calculating the Arrow-Pratt coefficients of absolute risk-aversion RA_U and RA_V for the two functions U and V, under the assumption that the density $f(x)$ is differentiable:

Corollary 3.1. *If $f(x)$ is differentiable and $f(x) > 0$, then*

$$RA_V(x) = -\frac{f'(x)}{f(x)} \quad \textit{and} \quad RA_U(x) = -\frac{2}{3}\frac{f'(x)}{f(x)}.$$

Assume, for example, that $X = [0, b],\ b > 0$, and alternatives are drawn from X according to a truncated exponential distribution with rate parameter $\lambda > 0$, so that $f(x) = (\lambda e^{-\lambda x})/(1 - e^{-\lambda b})$. The parameter λ measures how frequently choices involve small fitness levels. As λ grows, probability mass is shifted towards smaller alternatives. It now follows from Corollary 3.1 that $RA_V = \lambda$ and $RA_U = \frac{2}{3}\lambda$; both $U(x)$ and $V(x)$ exhibit constant absolute risk-aversion. Risk-aversion is, however, still decreasing in the sense that a decrease in λ, which corresponds to a shift of probability mass to larger payoffs, reduces risk-aversion.

Assume instead that $f(x)$ is single-peaked with peak in the interior of X. This seems to be a sensible description of both hunter-gatherers' environment and today's decision situations, where most opportunities involve middle-sized rather than extremely small or large payoffs.

Corollary 3.2. *If f is continuous and single-peaked with peak at $\hat{x} \in]a, b[$, then both $U(x)$ and $V(x)$ are S-shaped with inflection point $\hat{x}$.*

A utility function as described in the corollary resembles the value function used in prospect theory. Clearly, $\hat{x}$ can be interpreted as the decision maker's "reference point", because decisions involving alternatives close to $\hat{x}$ are most likely and the agent will be accustomed to decisions at this level. The advantage of the present approach is first that it gives an evolutionary rationale for S-shaped value functions in general[23], and second that it explains how an individual's reference point is determined. Even though the environment is stochastic and agents' payoffs will fluctuate over time, the reference point does not adjust to any newly experienced payoff, but should remain fixed as long as the density $f(x)$ remains the same. Adjustments of the reference point should therefore not be expected in response to random payoff realizations of exceptional size, but only to systematic changes of the environment, which in turn might be indicated by an accumulation of previously uncommon payoffs.

An additional feature of prospect theory is "loss aversion", captured by a downward kink of the value function at the reference point. In the present framework, this requires a downward jump of the fitness density at its peak, which may appear as a rather ad hoc assumption. There are, however, reasonable models in which exactly this discontinuity arises. Reconsider the discrete time model from Section 3.3, with just two time periods ($T = 1$). To capture the idea that instantaneous payoffs are usually smaller than later payoffs, for example due to natural growth, assume

[23] Rayo and Becker (2007a) show that their step function can become S-shaped if evolution cannot incorporate all relevant information into the happiness function.

that $g_0(v)$ is strictly decreasing while $g_1(v)$ is strictly increasing in v. It then immediately follows that the induced fitness density $f(x)$ is decreasing in x for $x > \delta$. If, in addition, $g_1(v)$ is increasing strongly enough, $f(x)$ will be increasing in x for $x \leq \delta$. The density $f(x)$ will then be single-peaked with a downward jump at the peak $\hat{x} = \delta$,[24] and the utility $U(0, v)$ exhibits loss aversion. In addition to discounting, there are other interpretations of the model that should be pointed at. An alternative $(0, v)$ could represent a project with payoff v that an agent can carry out by itself, while the project $(1, v)$ requires the help of a collaborator who receives a share $1 - \delta$ of the payoff. The assumption that large payoffs are more frequent in team projects than in individual projects is then again sensible. Similar conclusions can be derived from other models in which alternatives differ with respect to a binary characteristic.

Finally, assume that alternatives are drawn from a set Y and mapped into fitness by a function ψ, where Assumption 3.1 is again made for convenience. Also assume that ψ is twice differentiable and strictly concave. The elements of Y could, for example, represent bundles of food containing different amounts of nutriments, so that concavity of ψ amounts to the natural assumption of decreasing returns of nutrition with respect to fitness production. Concavity of ψ is clearly a potential reason for risk-aversion. Assume, for example, that $Y = X = [0, 1]$ and $g(y) = 1$, i.e. alternatives are drawn uniformly from Y, so that Proposition 3.1 implies

$$U''(y) = \frac{1}{3} c \, \psi'(y)^{-2/3} \psi''(y).$$

In this case, the curvature properties of $U(y)$ are directly inherited from ψ. This result is, however, driven by the assumption of a uniform distribution, which eliminates effects of environmental randomness. If $g(y)$ is single-peaked, $U(y)$ can be S-shaped even if the fitness function is concave. It follows from Proposition 3.1 that this will be the case whenever $g'(y)$ is sufficiently large and/or $g(y)$ is sufficiently small for small values of y. The inflection point will then always be left of the peak of g. Consider the following example, where $Y = X = [0, 1]$ and the fitness function is given by $\psi(y) = \sqrt{y}$. Assume that alternatives are drawn from Y according to a symmetric beta distribution with parameters $\alpha = \beta = 2$, which implies that

[24] Since g_1 is increasing, the jump discontinuity does actually exist.

the density $g(y) = 6y - 6y^2$ is single-peaked with peak at $\hat{y} = \frac{1}{2}$. It is now an easy exercise to show that $U''(y) > 0$ if $0 < y < \frac{3}{7}$ and $U''(y) < 0$ if $\frac{3}{7} < y < 1$. Hence utility is again S-shaped, with inflection point at $\frac{3}{7}$, which lies somewhat to the left of $\hat{y}$.

3.5 Conclusions

Under the assumption that human decisions are motivated by the pursuit of happiness, this chapter derives optimal hedonic utility functions as situation-specific tools for evolutionary success. If utility can only be perceived in discrete shades, different utility functions are differently well adapted to a choice situation. Evolution will select a function which is steep in regions where decisions have to be made frequently and errors are especially harmful. If these characteristics differ between choice situations, hedonic utility will adjust. Application of this insight yields evolutionary explanations for well-documented patterns of risk attitudes, for changing degrees of impatience, and for the coexistence of potentially conflicting utility functions.

The general model of hedonic adaption reveals that the slope of utility can be interpreted as the degree of attention devoted to the respective area. The central result then confirms the intuition that maximal fitness can be attained by allocating attention according to cost-benefit considerations. This economic argument might provide explanations for several behavioral patterns that present anomalies for the standard economic approach, among them what has been described as "mental accounting" (Thaler (1999)) or "choice bracketing" (Read et al. (1999)). Both theories are related to the multiple selves approach to time discounting (Shefrin and Thaler (1988)), and the present results might help to understand why different "accounts" or "brackets" exist and under which circumstances they become active.

3.A Appendix

3.A.1 Proof of Theorem 3.1

Given a density f, the expected fitness loss due to wrong decisions can be written as follows. Assume a first alternative $x \in [t_k, t_{k+1}]$ has been drawn from between two neighboring thresholds (or a boundary, respectively) that are located at positions t_k and t_{k+1}, where $t_k \leq t_{k+1}$. Then

$$L(x|t_k, t_{k+1}) = \int_{t_k}^{t_{k+1}} \frac{1}{2}|y - x| f(y) dy$$

is the expected loss conditional on x. The unconditional expected loss between t_k and t_{k+1} becomes

$$L(t_k, t_{k+1}) = \int_{t_k}^{t_{k+1}} L(x|t_k, t_{k+1}) f(x) dx.$$

The overall loss of a threshold allocation is obtained by adding this expression for all intervals between thresholds (and the boundaries).

Now consider the step density $\hat{f}_S$ as defined in section 3.2 and examine two neighboring thresholds at $t_k, t_{k+1} \in X_i$ for some $i \in \{1, ..., SD\}$. It follows that

$$L(t_k, t_{k+1}) = \frac{1}{6}(f_i)^2 (t_{k+1} - t_k)^3.$$

Consider first the problem of optimal threshold positions under the constraint that exactly N_i thresholds are allocated to step $i = 1, ..., SD$. Whenever $f_i = 0$ on some step i, $N_i = 0$ will clearly be optimal. All following arguments then apply unaltered by simply passing over this step. Hence for the moment assume $f_i > 0$ and $N_i \geq 1$ for all $i = 1, ..., SD$. Whenever $N_i \geq 3$, all thresholds in X_i must clearly be equidistant. This follows from observing that the distance between two thresholds enters the loss as a cubic term. Hence whenever $N_i \geq 2$, the thresholds span $N_i - 1$ intervals of length l_i in the interior of X_i, where the dependency of l_i on the whole profile $N_1, ..., N_{SD}$ is omitted for notational simplicity. There is one additional interval between a and the first threshold, and one additional interval between the largest threshold and b. Furthermore, for each $i = 1, ..., SD - 1$ there is one interval between the largest threshold in X_i and the smallest threshold in X_{i+1}.

Let $N_1(N), ..., N_{SD}(N)$ describe the optimal number of thresholds on each step if N thresholds are available altogether, which satisfies $\sum_{i=1}^{SD} N_i(N) = N$.[25] As $N \to \infty$, clearly $N_i(N) \to \infty$ for at least one step i, which implies $\lim_{N\to\infty} l_i = 0$. Assume that one interior threshold is removed from X_i, while all other thresholds remain unchanged. This increases the loss by $\frac{1}{6}(f_i)^2(2l_i)^3 - \frac{2}{6}(f_i)^2(l_i)^3 = (f_i)^2(l_i)^3$, which goes to zero as $N \to \infty$. This implies that the distance between *any two* neighboring thresholds (or the boundaries a or b respectively) has to go to zero as $N \to \infty$. If it did not for two thresholds t_k and t_{k+1}, relocating an interior threshold from X_i to $]t_k, t_{k+1}[$ would eventually (for large enough N) decrease the overall loss. Hence, $N_i(N) \to \infty$ as $N \to \infty$ for all $i = 1, ..., SD$. Furthermore, $\lim_{N\to\infty}(N_i(N) - 1)l_i = L(X_i)$.

Now consider a stronger necessary condition for optimality of $N_i(N)$, $i = 1, ..., SD$, where N is assumed to be large enough to imply $N_i(N) \geq 3$ for all $i = 1, ..., SD$. After taking one interior threshold out of step X_i, keep only the smallest and the largest threshold in X_i fixed, and rearrange the remaining thresholds in between to make them equidistant again. This increases the loss by

$$\frac{1}{6}(f_i)^2(l_i)^3(N_i - 1)^3 \left[\frac{1}{(N_i - 2)^2} - \frac{1}{(N_i - 1)^2}\right].$$

Similarly, keep the largest and smallest threshold in X_j, $j \neq i$, fixed, add the additional threshold in between, and rearrange to equidistant positions. This decreases the loss by

$$\frac{1}{6}(f_j)^2(l_j)^3(N_j - 1)^3 \left[\frac{1}{(N_j - 1)^2} - \frac{1}{(N_j)^2}\right].$$

The condition for this not to decrease the overall loss can be rearranged to

$$\left(\frac{N_j}{N_i - 1}\right)^2 \left(\frac{N_j - 1}{N_i - 2}\right)^2 \left(\frac{2N_i - 3}{2N_j - 1}\right) \geq \left(\frac{f_j}{f_i}\right)^2 \left[\frac{(N_j - 1)l_j}{(N_i - 1)l_i}\right]^3. \quad (3.1)$$

If the same argument is repeated for the relocation of a threshold from step j to step i, one obtains

$$\left(\frac{N_j - 1}{N_i}\right)^2 \left(\frac{N_j - 2}{N_i - 1}\right)^2 \left(\frac{2N_i - 1}{2N_j - 3}\right) \leq \left(\frac{f_j}{f_i}\right)^2 \left[\frac{(N_j - 1)l_j}{(N_i - 1)l_i}\right]^3. \quad (3.2)$$

[25]The dependency of $N_i, i = 1, ..., SD$, on N will sometimes be omitted for notational simplicity in the following.

As $N \to \infty$, the identical RHS of (3.1) and (3.2) converges to $(f_j/f_i)^2(L(X_j)/L(X_i))^3$ from the above considerations. Denote the LHS of (3.1) by $a_{ij}(N)$ and the LHS of (3.2) by $b_{ij}(N)$. Since $N_i, N_j \to \infty$ as $N \to \infty$, it follows that $\lim_{N\to\infty}(a_{ij}(N)/b_{ij}(N)) = 1$. It then follows by a straightforward argument that $\lim_{N\to\infty} a_{ij}(N) = \lim_{N\to\infty} b_{ij}(N) = (f_j/f_i)^2(L(X_j)/L(X_i))^3$ must hold, since otherwise either (3.1) or (3.2) would be violated for large N. Given existence of this limit, it also holds that $\lim_{N\to\infty} a_{ij}(N) = \lim_{N\to\infty} b_{ij}(N) = \lim_{N\to\infty}(N_j/N_i)^3$, so that the limit optimality condition becomes

$$\lim_{N\to\infty} \frac{N_j(N)}{N_i(N)} = \left(\frac{f_j}{f_i}\right)^{2/3} \left(\frac{L(X_j)}{L(X_i)}\right) \tag{3.3}$$

for all $i, j = 1, ..., SD$, $i \neq j$.

By fixing i and adding (3.3) for all $j = 1, ..., SD$, it follows that

$$\lim_{N\to\infty} \frac{N_i(N)}{N} = \frac{f_i^{2/3} L(X_i)}{\sum_{j=1}^{SD} f_j^{2/3} L(X_j)}$$

for all $i = 1, ..., SD$. This condition now also applies to steps where $f_i = 0$.

Now examine $U_{N,S}(x) := \theta_{N,S}(x)/N$. Denote by $\sigma_S(x)$ the number of the step on which x is located, i.e. $x \in X_{\sigma_S(x)}$. It now follows that

$$U_S(x) := \lim_{N\to\infty} \frac{\theta_{N,S}(x)}{N} = \frac{\sum_{i=1}^{\sigma_S(x)-1} f_i^{2/3} L(X_i) \quad + \quad \gamma(x,S) f_{\sigma_S(x)}^{2/3}}{\sum_{j=1}^{SD} f_j^{2/3} L(X_j)}$$

where $\gamma(x, S) = x - x_{\sigma_S(x)-1}$ is the distance between x and the lower end of the step on which it is located. Observe that $f_i^{2/3}$ is a value taken by the function $f(x)^{2/3}$ somewhere on X_i, because f_i is taken by $f(x)$ somewhere on X_i. Furthermore, $f(x)^{2/3}$ is Riemann-integrable since it is bounded and has finitely many discontinuities on $[a, b]$. We thus obtain

$$U(x) := \lim_{S\to\infty} U_S(x) = c \int_a^x f(y)^{2/3} dy,$$

where $c = (\int_a^b f(y)^{2/3} dy)^{-1}$ is a normalizing constant.

The result on $V(x)$ stated in Theorem 3.1 follows easily by repeating all previous steps for the probability of mistakes, where the probability of a mistake between two thresholds $t_k, t_{k+1} \in X_i$ is $P(t_k, t_{k+1}) = \frac{1}{2}(f_i)^2(t_{k+1} - t_k)^2$, analogously to $L(t_k, t_{k+1})$ above. Hence for this case, the technique used here yields the same result that Robson (2001a) obtained.

References of Chapter 3

AINSLIE, G. AND N. HASLAM (1992): "Hyperbolic Discounting," in *Choice Over Time*, ed. by G. Loewenstein and J. Elster, New York: Russell Sage Foundation.

BECHARA, A., H. DAMASIO, D. TRANEL, AND A. DAMASIO (1997): "Deciding Advantageously Before Knowing the Advantageous Strategy," *Science*, 275, 1293–1295.

BERGSTROM, T. (1997): "Storage for Good Times and Bad: Of Rats and Men," Mimeo.

COOPER, S. AND R. KAPLAN (1982): "Adaptive "Coin-flipping": A Decision-Theoretic Examination of Natural Selection for Random Individual Variation," *Journal of Theoretical Biology*, 94, 135–151.

CURRY, P. (2001): "Decision Making Under Uncertainty and the Evolution of Interdependent Preferences," *Journal of Economic Theory*, 98, 357–369.

DAMASIO, A. (1994): *Descartes' Error: Emotion, Reason, and the Human Brain*, New York: Avon.

DASGUPTA, P. AND E. MASKIN (2005): "Uncertainty and Hyperbolic Discounting," *American Economic Review*, 95, 1290–1299.

DEKEL, E., J. C. ELY, AND O. YILANKAYA (2007): "Evolution of Preferences," *Review of Economic Studies*, 74, 685–704.

ELSTER, J. (1985): "Weakness of Will and the Free-Rider Problem," *Economics and Philosophy*, 1, 231–265.

FREDERICK, S. AND G. LOEWENSTEIN (1999): "Hedonic Adaption," in *Well-Being: The Foundations of Hedonic Psychology*, ed. by D. E. Kahneman, D. and N. Schwarz, New York: Russell Sage Foundation, 302–329.

FREDERICK, S., G. LOEWENSTEIN, AND T. O'DONOGHUE (2002): "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40, 351–401.

GAZZANIGA, M., R. IVRY, AND G. MANGUN (1998): *Cognitive Neuroscience: The Biology of the Mind*, New York: Norton.

GUL, F. AND W. PESENDORFER (2001): "Temptation and Self-control," *Econometrica*, 69, 1403–1435.

HANSSON, I. AND C. STUART (1999): "Malthusian Selection of Preferences," *American Economic Review*, 80, 529–544.

KAHNEMAN, D. AND A. TVERSKY (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–291.

KOÇKESEN, L., E. A. OK, AND R. SETHI (2000a): "Evolution of Interdependent Preferences in Aggregative Games," *Games and Economic Behavior*, 31, 303–310.

——— (2000b): "The Strategic Advantage of Negatively Interdependent Preferences," *Journal of Economic Theory*, 92, 274–299.

KUPFERMANN, I., E. KANDEL, AND S. IVERSON (2000): "Motivational and Addictive States," in *Principles of Neural Science*, ed. by E. Kandel, J. Schwartz, and T. Jessell, New York: McGraw-Hill, 998–1013, 4 ed.

LAIBSON, D. (1999): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–477.

MCCLURE, S., D. LAIBSON, G. LOEWENSTEIN, AND J. COHEN (2004): "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, 306, 503–507.

POLLACK, R. (1968): "Consistent Planning," *Review of Economic Studies*, 35, 201–208.

RAYO, L. AND G. BECKER (2007a): "Evolutionary Efficiency and Happiness," *Journal of Political Economy*, 115, 302–337.

——— (2007b): "Habits, Peers, and Happiness: An Evolutionary Perspective," *American Economic Review, Papers and Proceedings*, 97, 487–491.

READ, D. (2001): "Intrapersonal Dilemmas," *Human Relations*, 54, 1093–1117.

READ, D., G. LOEWENSTEIN, AND M. RABIN (1999): "Choice Bracketing," *Journal of Risk and Uncertainty*, 19, 171–197.

ROBSON, A. (1996): "A Biological Basis for Expected and Non-Excpected Utility," *Journal of Economic Theory*, 68, 397–424.

——— (2001a): "The Biological Basis of Economic Behavior," *Journal of Economic Literature*, 39, 11–33.

——— (2001b): "Why Would Nature Give Individuals Utility Functions?" *Journal of Political Economy*, 109, 900–914.

ROBSON, A. AND H. KAPLAN (2007): "Why We Grow Large and then Grow Old: Economics, Biology and Mortality," Mimeo.

ROBSON, A. AND L. SAMUELSON (2007): "The Evolution of Intertemporal Preferences," *American Economic Review*, 97, 496–500.

SAMUELSON, L. (2004): "Information-Based Relative Consumption Effects," *Econometrica*, 72, 93–118.

SCHELLING, T. (1984): "Self-Command in Practice, in Policy, and in a Theory of Rational Choice," *American Economic Review*, 74, 1–11.

SHEFRIN, H. AND R. THALER (1988): "The Behavioral Life-Cycle Hypothesis," *Economic Enquiry*, 26, 609–643.

SOZOU, P. (1998): "On Hyperbolic Discounting and Uncertain Hazard Rates," *Proceedings of the Royal Society of London Series B*, 265, 2015–2020.

STROTZ, R. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165–180.

THALER, R. (1999): "Mental Accounting Matters," *Journal of Behavioral Decision Making*, 12, 183–206.

WINSTON, G. (1980): "Addiction and Backsliding: A Theory of Compulsive Consumption," *Journal of Economic Behavior and Organization*, 1, 295–324.

WRIGHT, S. (1932): "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution," *Proceedings of the VI International Congress of Genetics*, 1, 356–366.

CHAPTER 4

The Logit-Response Dynamics

4.1 Introduction and Literature

Models of learning in games typically start with the specification of a basic behavioral rule on the side of the players, e.g. myopic best reply, truncated fictitious play, or a variant of imitation. Since such basic dynamics exhibit a multiplicity of rest points (e.g., any Nash equilibrium is a rest point for a myopic best reply dynamics), it is necessary to perform a stability test.

Within the class of discrete-time, finite population models, one of the most successful paradigms in the literature performs this test by adding *noise* to the basic dynamics and studying the long-run outcomes as noise vanishes. Formally, the basic dynamics is a Markov chain[1] with multiple absorbing sets, which is made irreducible by the addition of noise. Probably the best-known example of this methodology is the *mistakes model*, essentially introduced by Kandori et al. (1993) (for an imitation rule), Young (1993), and Kandori and Rob (1995) (for myopic best reply). In this model, agents are assumed to have a certain probability (independent across agents and periods) of making mistakes, where a mistake is defined as choosing some strategy at random, with a full-support probability distribution. The mistake distribution is typically assumed to be uniform, although this is of no relevance. The important feature of the model is that the shape of this distribution is independent of the noise level.

The mistakes model has delivered important messages, ranging from the almost universal selection of risk-dominant equilibria (as opposed to Pareto efficient ones) in coordination games (Kandori et al. (1993), Young (1993)), to the dynamic relevance of "perfectly competitive" outcomes in aggregative games (Vega-Redondo (1997), Alós-Ferrer and Ania (2005)). One of the most attractive features of the mistakes model is that, thanks to a result due to Freidlin and Wentzell (1988), it is possible to provide a simple characterization of the set of long-run outcomes. These outcomes, called *stochastically stable states*, are those having positive probability in the limit invariant distribution as noise vanishes. The well-known characterization relies on minimizing the number of mutations associated to the transitions depicted in certain graphs (trees) defined on the state space.

[1]Throughout this chapter, the term "Markov chain" refers to a discrete-time Markov chain with stationary transition probabilities and finite state space.

This approach is not exempt of critiques. First and foremost, selection results are based on the number of mistakes necessary to destabilize a given state, but, in a sense, all mistakes are treated equally. For example, choosing a strategy which delivers *almost* a best response is a mistake at the same level as choosing a strategy which delivers payoffs far away from the optimum. Thus, the approach only relies on the (cardinal) payoffs of the game to a limited extent. Second, Bergin and Lipman (1996) observed that, if the distribution of mistakes is allowed to be state-dependent, the model can be twisted to select any pre-specified rest point. Thus, it becomes necessary to have a well-justified theory on the origin of mistakes. In van Damme and Weibull (2002), mistake probabilities are derived from a rationalistic model where agents can control mistakes at a cost, but it is prohibitive to completely eliminate them. Blume (2003) responds to the critique by characterizing the class of noise processes that lead to the selection of risk-dominant equilibria in coordination games.

The mistakes model, though, is not the only model where noise can be reasonably used as a selection device. One particular model which is immune to Bergin and Lipman's (1996) critique is the *logit response* dynamics, which was introduced in Blume (1993). In this dynamics, players adopt an action according to a full-support distribution of the logit form, which allocates larger probability to those actions which would deliver (myopically) larger payoffs. Thus, this dynamics is based on a specific theory about the origin of mistakes which in turn takes the magnitude of (suboptimal) payoffs fully into account. Noise is incorporated in the specification from the onset, and it governs directly the distribution of actions itself. As noise vanishes, that distribution concentrates on the best responses.

Due to the specification of noise, the logit-response dynamics is not a particular case of the mistakes model, and thus cannot benefit from the characterization of long-run outcomes mentioned above. Indeed, results for the logit-response dynamics are harder to obtain and are often restricted to particularly well-behaved classes of games. For example, binary action games (as in Blume (2003) or Maruta (2002)) give rise to a birth-death chain whose invariant distribution can be characterized directly. Further, as shown by Blume (1997), if the base game admits an exact potential, the process is reversible and again the invariant distribution can be characterized directly.

An additional criticism to the literature of learning in games, and the mistakes model in particular, is the sensitivity of the results to the specification of the dynamics, most notably revision opportunities (e.g. inertia). It is often the case that "who learns when" is as important as "who learns how".

In this chapter we develop a characterization of the long-run outcomes of the logit-response dynamics for arbitrary finite normal-form games. Furthermore, our result applies to a generalization of the original logit-response dynamics. In particular, we allow for an arbitrary specification of revision opportunities, encompassing e.g. independent revision opportunities (as in most versions of the mistakes model) and asynchronous learning (as in Blume's (1993) model). In order to obtain our results, we build on the analysis of Freidlin and Wentzell (1988) to characterize the invariant distribution of the logit-response dynamics for fixed noise levels, and use it to develop a characterization of the stochastically stable states.

In order to illustrate the method and its applicability, we proceed to study the convergence of the logit-response model for the various generalizations of the concept of potential game. Our method allows us to offer simple answers to several open questions. We find that, first, convergence to the set of Nash equilibria cannot be guaranteed for Monderer and Shapley's (1996) generalized ordinal potential games, but, second, convergence does obtain for Voorneveld's (2000) best-response potential games. We also show that the latter result is robust to the specification of revision opportunities under an additional condition which is satisfied both by independent inertia and asynchronous learning models. Third, we study the value of potential maximizers as an equilibrium refinement and find that the selection of potential maximizers (which obtains for exact potential games under Blume's (1993) asynchronous learning dynamics) fails two robustness tests. First, it fails even for exact potential games if revision opportunities do not fall into the asynchronous learning category. Second, it fails for any generalization of potential games even if revision opportunities are asynchronous.

The chapter is organized as follows. Section 4.2 reviews the logit-response dynamics and introduces our generalized dynamics. Section 4.3 presents our characterization of stochastically stable states, whose (technical) proof is relegated to the Appendix. Section 4.4 applies this characterization to the logit-response dynamics in best-response potential games. Section 4.5 discusses a generalization of our results. Section 4.6 concludes.

4.2 The Logit-Response Dynamics

4.2.1 The Logit Choice Function

Let $\Gamma = (I, (S_i, u_i)_{i \in I})$ be a finite normal-form game with player set $I = \{1, 2, ..., N\}$, strategy sets S_i and payoff functions u_i defined on the set of pure strategy profiles $S = S_1 \times ... \times S_N$. For a given player i, denote by $S_{-i} = \prod_{j \neq i} S_j$ the set of pure strategy profiles of i's opponents. Following convention, we denote $s = (s_i, s_{-i}) \in S$ and $u_i(s_i, s_{-i}) = u_i(s)$.

The game is played by boundedly rational players, who behave as myopic best repliers, but tremble in their decisions. Every period, some set of players is chosen to update their actions. We will further specify revision opportunities below.

When given the chance to revise, player i observes the actions s_{-i} of the opponents. The probability of choosing action s_i given the current profile s_{-i} is given by the *logit choice function*

$$p_i(s_i, s_{-i}) = \frac{e^{\beta u_i(s_i, s_{-i})}}{\sum_{s_i' \in S_i} e^{\beta u_i(s_i', s_{-i})}}, \tag{4.1}$$

where $0 < \beta < \infty$.

The scalar β can be interpreted as an inverse measure of the level of noise in players' decisions. As $\beta \to \infty$, the described rule converges to the myopic best reply rule. For any $0 < \beta < \infty$, players choose non-best replies with positive probability, but actions that yield smaller payoffs are chosen with smaller probability. The dynamic adjustment process defines an irreducible and aperiodic Markov chain $\{X_t^\beta\}_{t \in \mathbb{N}}$ on the state space S, with stationary transition probabilities $P_{ss'}^\beta = \Pr^\beta(s_t = s' | s_{t-1} = s)$ and (unique) invariant distribution μ^β.

For any specification of revision opportunities, we will refer to this dynamics as a *logit response dynamics*. Consider the particular case where exactly one player is randomly selected each period to revise his strategy,[2] and let $q_i > 0$ denote the probability that player i is selected. For this case, which we will refer to as *asynchronous learning*, the dynamics was first introduced by Blume (1993)[3] and has been further developed in e.g. Blume (1997, 2003), Young (1998), and Baron, Durieu, Haller, and Solal (2002a, 2002b). Taken as a behavioral rule, the underlying logit choice function (4.1) is rooted in the psychology literature (Thurstone (1927)). From the microeconomic point of view, it can be given a justification in terms of a random-utility model (see e.g. McKelvey and Palfrey (1995) for details). Hofbauer and Sandholm (2002, Section 2) observe that it is also the only choice function of the "quantal" form

$$C_i(u_i) = \frac{w(u_i(s_i, s_{-i}))}{\sum_{s_i' \in S_i} w(u_i(s_i', s_{-i}))}$$

with $w(\cdot)$ an increasing and differentiable function of the payoffs, which can be derived as the result of both a stochastic and a deterministic perturbation of the payoffs.[4] Thus, the logit-response dynamics exhibits solid decision-theoretic foundations.

4.2.2 Asynchronous Logit Response in Potential Games

The game Γ is a *potential game*[5] (Monderer and Shapley (1996)) if there exists a function $\rho : S \to \mathbb{R}$, called the potential, such that for each $i \in I$, $s_i, s_i' \in S_i$, $s_{-i} \in S_{-i}$

$$u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) = \rho(s_i, s_{-i}) - \rho(s_i', s_{-i}).$$

The global maximizers of the potential function ρ form a subset of the set of Nash equilibria of Γ. If Γ is a potential game, it follows that $u_i(s_i, s_{-i}) = \rho(s_i, s_{-i}) + \lambda(s_{-i})$, where $\lambda(s_{-i})$ is independent of s_i.[6] Thus (4.1) can be simplified to

$$p_i(s_i, s_{-i}) = \frac{e^{\beta u_i(s_i, s_{-i})}}{\sum_{s_i' \in S_i} e^{\beta u_i(s_i', s_{-i})}} = \frac{e^{\beta \rho(s_i, s_{-i})}}{\sum_{s_i' \in S_i} e^{\beta \rho(s_i', s_{-i})}}. \quad (4.2)$$

[2]This can be interpreted as a reduced form (technically, the *embedded chain*) of a continuous-time model where players receive revision opportunities according to "Poisson alarm clocks."

[3]Blume (1993) refers to this dynamics as log-linear response.

[4]Mattsson and Weibull (2002) and Baron, Durieu, Haller, and Solal (2002a, 2002b) show that logit-response arises in the framework of van Damme and Weibull (2002) when control costs adopt a specific functional form.

[5]Also called *partnership games*. See Hofbauer and Sigmund (1988).

[6]Fix a strategy $s_0 \in S_i$, and define $\lambda(s_{-i}) = u_i(s_0, s_{-i}) - \rho(s_0, s_{-i})$.

It is then straightforward to show (see Blume (1997)) that the invariant distribution of the logit-response dynamics adopts a *Gibbs-Boltzmann* form, i.e. the potential function becomes a potential for the stochastic process. The proof (which is included for completeness only) takes advantage of the fact that the reformulation (4.2) implies that the process is reversible.

Proposition 4.1. *Let Γ be a potential game with potential ρ. The invariant distribution of the logit-response dynamics with asynchronous learning is*

$$\mu^{\beta}(s) = \frac{e^{\beta\rho(s)}}{\sum_{s' \in S} e^{\beta\rho(s')}}.$$

Proof. It is enough to show that μ^{β} as given in the statement satisfies the *detailed balance condition*, i.e. $\mu^{\beta}(s)P^{\beta}_{ss'} = \mu^{\beta}(s')P^{\beta}_{s's}$ for all $s, s' \in S$. This is clearly fulfilled if $s = s'$, and also if s and s' differ in more than one coordinate, since $P^{\beta}_{ss'} = P^{\beta}_{s's} = 0$ in this case. Hence assume w.l.o.g. that s and s' differ exactly in coordinate i, that is $s_i \neq s'_i$ and $s_j = s'_j$ for all $j \neq i$. It follows that

$$\mu^{\beta}(s)P^{\beta}_{ss'} = \frac{e^{\beta\rho(s)}}{\sum_{s'' \in S} e^{\beta\rho(s'')}} \, q_i \, \frac{e^{\beta\rho(s'_i, s_{-i})}}{\sum_{s''_i \in S_i} e^{\beta\rho(s''_i, s_{-i})}} = \mu^{\beta}(s')P^{\beta}_{s's}.$$

where the last equality holds due to $s_{-i} = s'_{-i}$. □

As $\beta \to \infty$, the invariant distribution of the process converges to an invariant distribution of the best-reply dynamics. We say that a state $s \in S$ is *stochastically stable* if $\lim_{\beta\to\infty} \mu^{\beta}(s) > 0$. An immediate consequence of Proposition 4.1 is

Corollary 4.1. *Let Γ be a potential game. The set of stochastically stable states of the logit-response dynamics with asynchronous learning is equal to the set of maximizers of ρ.*

This Corollary provides of course a readily applicable result.[7] In our view, it is also important for two additional reasons. First, it is a *convergence* result. The asynchronous logit response dynamics always converges to the set of Nash equilibria in the class of exact potential games. Second, it is a *selection* result. In particular, the logit-response dynamics provides support for treating the set of potential maximizers as an equilibrium refinement for potential games.

[7]For example, Sandholm (2007) relies on this result to build a model of evolutionary implementation.

The latter finding has motivated a large part of the literature of learning in games in recent years, and indeed the selection of potential maximizers has become a test of the reasonability of a learning dynamics.[8] It is therefore important to know how robust both parts of Corollary 4.1 are. That is, we pose the question of whether the convergence to Nash equilibria and the selection of potential maximizers extend to more general classes of games and dynamics.

Proposition 4.1 (and hence Corollary 4.1), however, rely on the knife-edge technical fact that the exact potential of the game allows to identify the invariant distribution of the stochastic process for positive noise level. Clearly, the proof cannot be generalized any further. In the next section, we develop a framework which will allow us to provide exact results for both more general games and more general dynamics.

4.2.3 *Revision Processes and a Generalized Dynamics*

The existing results for the logit-response rule (as e.g. Corollary 4.1) rely on the asynchronicity assumption to establish the convenient Gibbs-Boltzmann form for the invariant distribution. Here we will consider a more general approach allowing for arbitrary specification of updating opportunities. We illustrate this by considering a general class of revision processes. The motivation for the generalization is as follows. In our view, a learning dynamics in games is made of a behavioral rule and a specification of revision opportunities (i.e. the speed of the dynamics). Thus, it is important to know which results are due to the behavioral rule and which ones hinge on the exact specification of the revision process. Studying general revision processes for a given dynamics therefore becomes an important robustness check.

Definition 4.1. *A* revision process *is a probability measure q on the set of subsets of I, $\mathcal{P}(I)$, such that*

$$\forall\, i \in I\ \exists J \subseteq I \text{ such that } i \in J \text{ and } q_J > 0 \tag{4.3}$$

where, for each $J \subseteq I$, $q_J = q(J)$ is interpreted as the probability that exactly *players in J receive revision opportunities (independently across periods).*

Condition (4.3) merely specifies that every player has some probability of being able to revise in some situation. No further restriction is placed on the revision

[8] See e.g. Hofbauer and Sorger (1999).

process, which allows for a wide range of models to be considered. We list now three leading examples.

Let $\mathcal{R}^q = \{J \subseteq I | q_J > 0\}$ denote the set of *revising sets*, i.e. sets of players which might obtain revision opportunity (as a whole) with positive probability. If $\mathcal{R}^q = \{\{i\} | i \in I\}$, we say that the dynamics exhibits *asynchronous learning* and write $q_i = q_{\{i\}}$. As commented above, this includes the asynchronous logit-response dynamics of Blume (1993).

If $\mathcal{R}^q = \mathcal{P}(I)$, we speak of *independent learning*. That is, every subset of players has positive probability of being able to revise. For example, a standard version of the mistakes model (see e.g. Sandholm (1998)) is a particular case which postulates *independent inertia*, i.e. each player revises with a fixed, independent probability $0 < p < 1$. Thus $q_J = p^{|J|}(1-p)^{N-|J|} > 0$ for each subset J.

We can also consider a model of *instantaneous learning*, where all players receive revision opportunities every period, i.e. $\mathcal{R}^q = \{I\}$. Other examples could include specific correlation in revision opportunities among certain groups of players,[9] or bounds to the number of players revising each period.

Fix a revision process q. For any two strategy profiles $s, s' \in S$, let $R_{s,s'} = \{J \in \mathcal{R}^q | s'_k = s_k \forall\, k \notin J\}$ be the set of revising sets potentially leading from s to s'. Note that from a given $s \in S$, different alternative revising sets might give rise to the same transition, because players selected to revise might stay with their previous action. However, under asynchronous learning $|R_{s,s'}| \leq 1$ for all $s \neq s'$. We say that a transition from s to s' is *feasible* if $R_{s,s'} \neq \emptyset$.

The logit-response dynamics with revision process q is a Markov chain on the state space S with stationary transition probabilities given by

$$P_{s,s'} = \sum_{J \in R_{s,s'}} q_J \prod_{j \in J} \frac{e^{\beta \cdot u_j(s'_j, s_{-j})}}{\sum_{s''_j \in S_j} e^{\beta \cdot u_j(s''_j, s_{-j})}}.$$

Define $U_J(s', s) = \sum_{j \in J} u_j(s'_j, s_{-j})$. Let $R_s^J = \{s' \in S | s'_k = s_k \forall\, k \notin J\}$ be the set of states potentially reached from s when the revising set is J. We can then rewrite

[9] Since we do not restrict attention to symmetric games, this possibility might be of independent interest, e.g. for buyers-sellers models as in Alós-Ferrer and Kirchsteiger (2007).

the transition probabilities as

$$P_{s,s'} = \sum_{J \in R_{s,s'}} q_J \frac{e^{\beta \cdot U_J(s',s)}}{\sum_{s'' \in R_s^J} e^{\beta \cdot U_J(s'',s)}}. \tag{4.4}$$

4.3 Stochastic Stability

Given a general revision process q, the logit-response dynamics is in general not a birth-death chain. Even if this were the case (e.g. under asynchronicity in binary action games), unless the game is an (exact) potential game an exact characterization of the invariant distribution was until now not available. We now develop a characterization of the set of stochastically stable states of the logit-response dynamics which relies precisely on such a characterization of the invariant distribution.

Given a state s, define an *s-tree* to be a directed graph T such that there exists a unique path from any state $s' \in S$ to s. The key concept for our characterization is as follows:

Definition 4.2. *A* revision s-tree *is a pair* (T, γ) *where*

(i) T *is an s-tree,*

(ii) $(s, s') \in T$ *only if* $R_{s,s'} \neq \emptyset$ *(only feasible transitions are allowed), and*

(iii) $\gamma : T \to \mathcal{P}(I)$ *is such that* $\gamma(s, s') \in R_{s,s'}$ *for all* $(s, s') \in T$.

Thus, there are two differences between a revision tree and a tree as used in the characterization for the mistakes model. First, in a revision s-tree, edges corresponding to unfeasible transitions are not allowed.[10] Second, in a revision s-tree (T, γ), γ labels each edge of T with a revising set which makes the corresponding transition potentially feasible.

Remark 4.1. Suppose that a revision process satisfies that for all $s, s' \in S$, $s \neq s'$, either $\mathcal{R}_{s,s'} = \emptyset$ or $|\mathcal{R}_{s,s'}| = 1$. This is e.g. true for asynchronous learning and instantaneous learning. Then, for each link (s, s') in a revision tree there exists exactly one revising set making the transition from s to s' feasible. In other words, given a tree T using feasible transitions only, there exists a unique mapping γ such that (T, γ) is a revision tree.

[10]Thus, actually the concept of revision tree depends on the revision process q. We drop this dependency for notational simplicity.

4.3.1 A Characterization

Let $\mathcal{T}(s)$ denote the set of revision s-trees. The *waste* of a revision tree $(T, \gamma) \in \mathcal{T}(s)$ is defined as

$$W(T,\gamma) = \sum_{(s,s')\in T} \left(\max_{s''\in S} U_{\gamma(s,s')}(s'', s) \right) - U_{\gamma(s,s')}(s', s).$$

or, equivalently,

$$W(T,\gamma) = \sum_{(s,s')\in T} \sum_{j\in\gamma(s,s')} \left(\max_{s''_j \in S_j} u_j(s''_j, s_{-j}) - u_j(s'_j, s_{-j}) \right).$$

In words, the waste of a revision tree adds all the individual (ex-ante, myopic) payoff wastes generated across the transitions depicted in the tree, relative to the payoffs that could have been reached by adopting best responses. Obviously, a transition generates zero waste in this sum if and only if it involves only best responses.

Intuitively, the waste of a revision tree is an inverse measure for its likelihood in the logit-response dynamics. It is analogous to the concept of costs in the mistakes model, with the obvious difference that wastes are real numbers, rather than natural ones (number of mistakes).[11] The *stochastic potential* of a given state is obtained by minimizing waste across revision trees rooted in that state.

Definition 4.3. *The* stochastic potential *of a state s is*

$$W(s) = \min_{(T,\gamma)\in\mathcal{T}(s)} W(T,\gamma).$$

As mentioned above, a state is stochastically stable if it has positive probability in the limit invariant distribution of a noisy process as noise vanishes (in our case, when $\beta \to \infty$). Our characterization of stochastically stable states is as follows.

Theorem 4.1. *Consider the logit-response dynamics (with any revision process). A state is stochastically stable if and only if it minimizes $W(s)$ among all states.*

[11] An alternative name for the waste would be *regret*. We prefer to avoid this name for two reasons. First, there is a growing game-theoretic literature where players choose actions according to their associated regret (see e.g. Hart and Mas-Colell (2000)). For us, the waste is rather a technical device and not an objective target. Second, except in the case of asynchronous learning, the waste of a player's choice is only potential regret, since the corresponding payoff will not actually be experienced due to other players simultaneously updating their choices.

Proof. See Appendix, Section 4.A.1. □

The proof of Theorem 4.1 is itself based on an exact characterization of the invariant distribution for finite β (Lemma 4.1 in the Appendix). Theorem 4.1 yields a "tree-surgery" technique for the characterization of stochastically stable states of the logit-response dynamics, for arbitrary finite normal-form games and with arbitrary revision processes. It is analogous to the statement that stochastically stable states are those having trees involving a minimal number of mistakes in the mistakes model. In our framework, the number of mistakes is replaced by the sum of payoff losses from transitions which are not possible in the limit as $\beta \to \infty$.

This result makes it possible to focus on minimal waste revision trees to examine stochastic stability. If the set of revising sets is a singleton for every possible transition, as is the case e.g. for asynchronous and instantaneous learning, there is exactly one revision tree per tree involving feasible transitions only, and thus we can directly examine minimal waste trees.

4.3.2 *A Radius-Coradius Result*

One of the most powerful results for the actual analysis of models based on the mistakes formulation is the Radius-Coradius theorem due to Ellison (2000). In order to support our tractability claim for the logit model, we now prove a result analogous to Ellison's (2000) Radius-Coradius Theorem in our framework. A directed graph P on S is a *path* if there exists a finite, repetition-free sequence $(s_0, s_1, ..., s_n)$ of states in S with $n = |P|$, such that $(s_m, s_{m+1}) \in P$ and $R_{s_m, s_{m+1}} \neq \emptyset$ for all $m = 0, ..., n-1$. The state s_0 is the *initial* point of the path, the state s_n is the *terminal* point. Since the logit response dynamics is irreducible for any revision process, the set of paths between any two given states is nonempty.

Note that a path as described above is an s_n-tree on the subset of states $\{s_0, ..., s_n\}$ and thus a revision path (P, γ) can be defined as a revision tree, where P is a path. Denote the set of all revision paths with initial point s and terminal point s' by $\mathcal{P}(s, s')$. The waste $W(P, \gamma)$ of a revision path $(P, \gamma) \in \mathcal{P}(s, s')$ is simply its waste as a revision tree.

The *basin of attraction*[12] of a state s, $B(s) \subseteq S$ is the set of all states s' such that there exists a revision path $(P, \gamma) \in \mathcal{P}(s', s)$ with $W(P, \gamma) = 0$. The *limit set* of state s is the set of states which are connected back-and-forth with s at zero waste, i.e. $L(s) = \{s' \in s | s' \in B(s) \text{ and } s \in B(s')\}$.

The Radius of a state s is defined as

$$R(s) = \min\{W(P, \gamma) | s' \notin B(s), (P, \gamma) \in \mathcal{P}(s, s')\}$$

and is a measure of how easy it is to leave state s. Since the waste is based on payoff differences and not number of mistakes, it takes into account not only the size but also the "depth" of the basin of attraction. The Coradius of s is given by

$$CR(s) = \max_{s' \notin B(s)} \min\{W(P, \gamma) | s'' \in B(s), (P, \gamma) \in \mathcal{P}(s', s'')\}$$

and is a measure of how hard it is to reach s.

Proposition 4.2. *Suppose a state $s \in S$ is such that $R(s) > CR(s)$. Then, the stochastically stable states are exactly those in $L(s)$.*

Proof. Let $s^* \in S$, $s^* \notin B(s)$. Let $(T^*, \gamma^*) \in \mathcal{T}(s^*)$ solve $\min_{(T,\gamma) \in \mathcal{T}(s^*)} W(T, \gamma)$.

Consider the tree T^* and the complete path from s to s^* in this tree. Since $s^* \notin B(s)$, this path eventually leaves the basin of attraction of s. Let s_1 be the first state in this path which is *not* in $B(s)$. Delete the part of the path from s to s_1. For all states but s that have become disconnected, the fact that they are in $B(s)$ allows to connect them to s (adding the corresponding transitions) with waste zero. If this creates any duplicated edges in the graph, delete the duplicate, but keep only the revising set which ensures waste zero. This saves a waste weakly larger than $R(s)$ (by definition of Radius).

Add to the revision tree a revision path $(P, \gamma) \in \mathcal{P}(s^*, s)$ which solves $\min\{W(P, \gamma) | s'' \in B(s), (P, \gamma) \in \mathcal{P}(s^*, s'')\}$. Delete any duplicated transitions created when adding (P, γ), keeping the revising sets in γ. This increases the waste by weakly less than $CR(s)$ (by definition of Coradius).

After these two operations we have constructed a new revision tree, rooted in s. If $CR(s) < R(s)$, the total waste has been strictly reduced. It follows that the stochastic potential of s is strictly smaller than the stochastic potential of any s^*

[12]There is a subtle difference between our result and Ellison's (2000). Ellison defines the basin of attraction of a state s as the set of states from which the (unperturbed) dynamics will eventually lead to s with probability one, whereas we define the basin of attraction of s as the set of states such that the unperturbed dynamics (i.e. that involving zero-waste transitions only) leads to s with positive probability.

not in the basin of attraction of s, thus the latter can not be stochastically stable by Theorem 4.1.

Consider now a state $s^* \in B(s)$ such that $s \notin B(s^*)$. Consider a minimal-waste s^*-revision tree. Since $s \notin B(s^*)$, in the path connecting s to s^* contained in this tree there exists some transition, say from s_1 to s_2, which causes strictly positive waste. Delete it. Since $s^* \in B(s)$, there exists a zero-waste revision path from s^* to s. Add this path to the revision tree, deleting duplicated transitions. The result is an s_1-revision tree with strictly smaller waste, thus again by Theorem 4.1, s^* cannot be stochastically stable.

Last, consider any state $s^* \in L(s)$, $s^* \neq s$. Clearly, minimal waste revision trees for both states must have the same waste. In summary, no state out of $L(s)$ can be stochastically stable, but all states in $L(s)$ have the same stochastic potential. Since there are finitely many states, there must exist states with minimum stochastic potential and the conclusion follows. □

Following Ellison (2000), it is possible to extend this result in two ways. The first would allow to apply the analysis to sets of states rather than a single state. The second would deal with the concept of "modified coradius", which subtracts the radius of intermediate states when computing the coradius, thus providing a more involved but stronger result.

4.4 Learning in Best-Response Potential Games

In this Section, we illustrate the use of our characterization and provide definite answers to the questions we posed above, that is, to which extent are the findings of convergence to Nash equilibria and selection of potential maximizers robust. To check robustness with respect to the dynamics, we will consider arbitrary revision processes as discussed above. To check robustness with respect to the class of games, we will consider the various generalizations of the concept of potential game.

4.4.1 Generalized Potential Games

As mentioned in Section 4.2.2, a finite normal form game $\Gamma = (I, (S_i, u_i)_{i \in I})$ is an (exact) *potential game* if there exists a function $\rho : S \mapsto \mathbb{R}$ (the potential) such that

$$u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) = \rho(s_i, s_{-i}) - \rho(s_i', s_{-i}) \quad \text{(P)}$$

for all $i \in I$, $s_i, s_i' \in S_i$, and $s_{-i} \in S_{-i}$. The set of potential maximizers has been shown to be an appealing equilibrium refinement for this class of games. However, it can also be argued that the class of potential games is relatively narrow. Monderer and Shapley (1996) generalized this class as follows. Γ is a *weighted potential game* if (P) is replaced by $u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i}) = w_i\left(\rho(s_i, s_{-i}) - \rho(s'_i, s_{-i})\right)$ for fixed weights $w_i > 0$, $i \in I$. Further, Γ is an *ordinal potential game* if (P) is replaced by the property that $u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})$ and $\rho(s_i, s_{-i}) - \rho(s'_i, s_{-i})$ have the same sign. Last, Γ is a *generalized ordinal potential game* if (P) is replaced by the property that $u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i}) > 0$ implies that $\rho(s_i, s_{-i}) - \rho(s'_i, s_{-i}) > 0$.

The appeal of generalized ordinal potential games rests on the following characterization. A finite game is generalized ordinal potential if and only if it has the *Finite Improvement Property*, that is, if any path of states generated through unilateral deviations involving strict improvements is necessarily finite.

Obviously, every potential game is a weighted potential game, every weighted potential game is an ordinal potential game, and every ordinal potential game is generalized ordinal potential. Voorneveld (2000) has provided a different generalization of the class of ordinal potential games, and has shown that it is neither included in nor includes the class of generalized ordinal potential games. The game Γ is a *best-response potential game* if there exists a function $\rho^{BR} : S \to \mathbb{R}$ such that $\forall i \in I$, and $s_{-i} \in S_{-i}$,

$$\arg\max_{s_i \in S_i} u_i(s_i, s_{-i}) = \arg\max_{s_i \in S_i} \rho^{BR}(s_i, s_{-i}).$$

Best-response potential games admit a characterization as follows (see Voorneveld (2000, Theorem 3.2)). A normal form game with finitely many players and countable strategy sets is a best-response potential game if and only if any path of states generated through unilateral best responses, and containing at least one strict improvement, is non-cyclic.

4.4.2 *A Convergence Result*

We turn now to the question of convergence to Nash equilibria.[13] Theorem 4.1 allows the following first, immediate observation. In generalized ordinal poten-

[13] Hofbauer and Sandholm (2002) use stochastic approximation techniques to study convergence of closely-related dynamics to the set of Nash equilibria in potential and supermodular games. The strategy is taking the limit as the population size grows to infinity and the time interval goes to zero, and approximating the paths of the dynamics through a differential equation. In contrast, we study convergence directly on the finite, fixed-population-size, discrete-time dynamics. Baron et al. (2002a) have established convergence of the asynchronous logit-response dynamics to *partial* Nash configurations, i.e. strategy profiles where at least one player is choosing a best response.

tial games, convergence to Nash equilibria is not guaranteed, even under asynchronous learning. In other words, non-Nash states can be stochastically stable. To see this, consider the following example.

Example 4.1. Consider asynchronous learning. The following 2×2 game (left-hand-side table) is Example 4.1.(a) in Voorneveld (2000).

	a	b
a	0,0	0,1
b	0,1	1,0

Payoff Table

	a	b
a	0	1
b	3	2

G.O. Potential

The only pure-strategy Nash equilibrium is (b, a). This game has a generalized ordinal potential ρ given by the right-hand-side table. However, the game exhibits a best-response cycle, and hence is not a best-response potential game. This best-response cycle contains the links (aa, ab), (ab, bb), (bb, ba), and (ba, aa). Since each of these transitions is a best response for the updating player, we can construct a zero-waste revision tree for all four states. In conclusion, all four states are stochastically stable, even though only one of them is a Nash equilibrium.[14]

This example shows that non-Nash states can be stochastically stable in generalized ordinal potential games under the logit-response dynamics with asynchronous learning. Thus, the next question of interest is *when* does convergence to Nash equilibria obtain. The example above shows that the answer is negative for the class of generalized ordinal potential games.

As an application of Theorem 4.1, though, we can answer this question in the affirmative for best-response potential games, and hence ordinal potential games. We will also simultaneously perform a robustness check of the convergence result to variations in the way players are chosen to update strategies. Say that a revision process is *regular* if $q_{\{i\}} > 0$ for all $i \in I$. Both standard revision processes in the learning literature, asynchronous learning and independent learning, are clearly regular.

[14]In particular, state ab is among the stochastically stable states, even though the set of Nash equilibria is $\{[(p, 1-p), a] \,|0 \leq p \leq \frac{1}{2}\}$, i.e. player 2 never plays strategy b in an equilibrium.

Theorem 4.2. *If* Γ *is a finite best-response potential game, the set of stochastically stable states of the logit-response dynamics with any regular revision process is contained in the set of Nash equilibria.*

Proof. Fix a state $s^0 \in S$ which is not a Nash equilibrium of Γ, and hence there exists a coordinate $i \in I$ such that $\max_{s_i \in S_i} u_i(s_i, s^0_{-i}) > u_i(s^0)$. Consider any revision tree $(T^0, \gamma^0) \in \mathcal{T}(s^0)$ with associated waste $W(T^0, \gamma^0)$. Construct a revision tree (T^1, γ^1) from (T^0, γ^0) as follows. Let $s^1 = (s^1_i, s^0_{-i})$ where $s^1_i \in \arg\max_{s_i \in S_i} u_i(s_i, s^0_{-i})$. Add the link (s^0, s^1) with revising set $\{i\}$ (which is possible by regularity of the revision process) and delete the link (s^1, s^2) leaving s^1 in T^0. The new graph is a tree $T^1 \in \mathcal{T}(s^1)$. The additional transition from s^0 to s^1 causes no waste by definition of s^1. If the contribution of the deleted link (s^1, s^2) to $W(T^0, \gamma^0)$ was positive, $W(T^1, \gamma^1) < W(T^0, \gamma^0)$ holds.[15]

If the contribution was zero and thus $W(T^1, \gamma^1) = W(T^0, \gamma^0)$, proceed as follows. Add a link $(s^1, \hat{s}^2)$ with revising set $\{i\}$ where $\hat{s}^2 = (s^2_i, s^1_{-i})$ for some $i \in \gamma^0(s^1, s^2)$ such that $s^2_i \neq s^1_i$. This causes zero waste because (s^1, s^2) caused zero waste. Delete the link $(\hat{s}^2, s^3)$ leaving $\hat{s}^2$ in T^1. The new (labelled) graph is a revision-tree $(T^2, \gamma^2) \in \mathcal{T}(\hat{s}^2)$, with zero waste for the link $(s^1, \hat{s}^2)$.

Iterate the described procedure until deletion of a positive waste link $(\hat{s}^n, s^{n+1})$ occurs, i.e. move along a *best-response compatible* path of states. Since Γ is a best-response potential game, any such path, which started with a strict improvement for a player, is non-cyclic (Voorneveld (2000, Theorem 3.2)), such that iteration actually ends with a revision tree $(T^n, \gamma^n) \in \mathcal{T}(\hat{s}^n)$ where $\hat{s}^n \neq s^0$ and $W(T^n, \gamma^n) < W(T^0, \gamma^0)$. Hence no $(T, \gamma) \in \mathcal{T}(s^0)$ can have minimum waste and, by Theorem 4.1, s^0 is not stochastically stable. □

This result generalizes both the class of potential games and the class of logit-response dynamics for which convergence to Nash equilibria obtains. The proof relies crucially on the characterization of finite best-response potential games, that is, the property that any path of states generated through unilateral best responses, containing at least one strict improvement, is non-cyclic. This property is not necessarily fulfilled by generalized ordinal potential games (e.g. it fails in Example 4.1.)

The assumption of regularity of the revision process is necessary for the result to obtain. To see this, consider the case of instantaneous learning, where every player receives revision opportunity with probability one. In this case, convergence to Nash equilibria can fail even for exact potential games.

[15] We use the term "waste of a link" as a shortcut for "waste of the revision tree formed by a single link and the chosen revising set".

Example 4.2. The following 2×2 game is symmetric, and hence an exact potential game (and also a best-response potential game). It has two strict Nash equilibria, (a, a) and (b, b).

	a	b
a	1,1	0,0
b	0,0	1,1

Payoff Table

	a	b
a	1	0
b	0	1

Exact Potential

Under asynchronous learning, both Nash equilibria are stochastically stable, since they both maximize the potential. With our approach it is easy to verify that the same holds under independent learning. Now consider instantaneous learning. Once a Nash state is reached, a waste of 1 is required to leave it, i.e. one of the updating players needs to make a mistake to move to either (a, b) or (b, a).[16] Once the process reaches either (a, b) or (b, a), it alternates between these two states if nobody makes a mistake. Leaving this cycle again causes a waste of 1. Hence the stochastic potential of all states is 2, and they are all stochastically stable. That is, convergence to Nash equilibria might fail even for exact potential games.

4.4.3 *The Irrelevance of Potential Maximizers*

This leads us to the second question of interest, namely whether potential maximizers are selected by the logit response dynamics in general. The following example illustrates that states which globally maximize the potential function of a weighted potential game will not generally be stochastically stable. Thus, although all stochastically stable states of the logit-response dynamics are Nash equilibria for best-response potential games, stochastic stability does not support the use of potential maximizers as an equilibrium refinement for any generalization of potential games, even with asynchronous learning.

Example 4.3. Let Γ be an asymmetric, pure-coordination, 2×2 game with strategy sets $S_1 = S_2 = \{a, b\}$ and payoffs as given in the following (left-hand-side) table:

[16]Moving directly from one of the Nash states to the other causes a waste of 2, because both players must make a mistake.

	a	b
a	2,2	0,0
b	0,0	10,1

Payoff Table

	a	b
a	2	-6
b	0	4

Weighted Potential

This game has a weighted potential ρ given by the right-hand-side table and weights $w_1 = 1$ and $w_2 = 1/4$. The equilibrium (b, b) is the (unique) potential maximizer.

Consider asynchronous learning. It is straightforward to construct the minimum-waste trees. Note that, since the game is a strict coordination game, states (a, b) and (b, a) can be connected to either of the pure Nash equilibria at zero waste. Thus the minimum waste of a (b, b) tree is equal to the minimum waste necessary to leave (a, a), and vice versa. The waste of the link $(a, a) \mapsto (b, a)$ is $w_1 \cdot (2 - 0) = 2$; the waste of the link $(a, a) \mapsto (a, b)$ is $w_2 \cdot (2 - (-6)) = 2$. Thus the stochastic potential of (b, b) is 2. Consider now state (b, b). The waste of the link $(b, b) \mapsto (a, b)$ is $w_1 \cdot (4 - (-6)) = 10$; the waste of the link $(b, b) \mapsto (b, a)$ is $w_2 \cdot (4 - 0) = 1$. Hence the stochastic potential of (a, a) is 1 and we conclude that (a, a) is stochastically stable, despite not maximizing ρ. This result can also be derived using the Radius-Coradius Theorem. Obviously, $R(a, a) = CR(b, b) = 2$ and $CR(a, a) = R(b, b) = 1$, implying that (a, a) is stochastically stable.

This example shows that the selection of potential maximizers for the asynchronous logit-response is not robust even to slight generalizations of the class of potential games. Now we consider whether the result is robust to generalizations of the class of revision processes.

There are two major differences between the asynchronous-learning case and, say, a process with independent learning. First, the set of revision trees for each state grows, since transitions between any two states become possible. Second, each transition in which not all N players change their action becomes possible via more than one revising set.

Concerning stochastic stability, though, this second issue raises no difficulties. Consider the link (s, s') where the players in J change their action, and assume that a revising set $J' \supset J$ is selected for this transition. It is easy to see that the corresponding waste can only be larger than if the revising set J was selected instead, because sticking to their action might be a non-best response for players in $J' \setminus J$. Hence, when computing the stochastic potential of a state, we can restrict attention to selections for trees that pick the most "parsimonious" revising sets, which prescribes a unique selection for each tree.[17]

The larger set of trees can, however, substantially change other results. We proceed to show that, under independent learning, potential maximizers may fail to be selected even in exact potential games. Thus the result of Corollary 4.1 is not robust to changes in the specification of revision opportunities either.

Example 4.4. Consider a $3 \times 3 \times 2$-game with exact potential as given below. Player 1 chooses rows, player 2 chooses columns, and player 3 chooses tables. The payoffs of pure-strategy Nash equilibria are marked by an asterisk.

g

	d	e	f
a	10*	6	0
b	6	0	0
c	0	0	9*

h

	d	e	f
a	0	0	0
b	0	1*	1*
c	0	1*	1

Under asynchronous learning, the potential maximizing state (a, d, g) is stochastically stable by Corollary 4.1. Consider independent learning instead. The basin of attraction of state (c, f, g) contains all states except (a, d, g), (a, e, g) and (b, d, g). Any minimal waste path from (c, f, g) to one of these states, for example the path $((c, f, g), (a, f, g), (a, d, g))$ or $((c, f, g), (c, f, h), (c, d, h), (b, d, g))$ is associated with a waste of 9, such that $R(c, f, g) = 9$. The transition $((a, d, g), (b, e, g))$, though, has an associated waste of only 8 when players 1 and 2 switch simultaneously. The states (a, e, g) and (b, d, g) can be connected to $B(c, f, g)$ at an even lower waste, such that $CR(c, f, g) = 8$. Proposition 4.2 then implies that (c, f, g) is stochastically stable, despite the fact that it does not maximize the exact potential.

[17]Essentially, that is the reason why Theorem 4.2 holds for any regular revision process.

4.5 Generalizations and Extensions

Although we have focused on the logit-response dynamics, our approach to stochastic stability is susceptible of generalization to a wider class of learning processes. In this Section, we briefly report on this generalization.

Consider a Markov chain $\{X_t\}_{t\in\mathbb{N}}$ on a finite state space Ω. Denote the stationary transition probabilities by $P_{\omega,\omega'} = Pr(X_t = \omega'|X_{t-1} = \omega)$. A *transition mechanism* from state ω is a mapping $Q : \Omega \to \mathbb{R}_+$. The interpretation is that from a given state ω, there might be different, alternative processes giving rise to a transition to other states. Conditional on the transition mechanism Q being selected, a state $\omega' \in \Omega$ will be reached from ω with probability

$$Q(\omega)/\sum_{\omega''\in\Omega} Q(\omega'').$$

Denote by M_ω the set of transition mechanisms available at ω, and let $M = \bigcup_{\omega\in\Omega} M_\omega$. Note that the sets M_ω need not be pairwise disjoint, so that a transition mechanism might be available at several or even all states (e.g. a random mutation). Further, let $M_{\omega,\omega'} = \{Q \in M_\omega|Q(\omega') > 0\}$, i.e. the set of mechanisms which are available at ω and may lead to ω'.

Definition 4.4. *Let X_t be a Markov chain on the finite state space Ω. A* decomposition *of X_t is a tuple $(M_\omega, q_\omega)_{\omega\in\Omega}$ such that, for each $\omega \in \Omega$,*

(i) M_ω is a nonempty, finite set of transition mechanisms,

(ii) $q_\omega \in \Delta M_\omega$ is a full-support probability measure on M_ω, and

(iii) for each $\omega' \in \Omega$,

$$P_{\omega,\omega'} = \sum_{Q\in M_{\omega,\omega'}} q_\omega(Q)\frac{Q(\omega')}{\sum_{\omega''\in\Omega} Q(\omega'')}.$$

Obviously, any finite Markov chain admits a trivial (and not very useful) decomposition with $M_\omega = \{Q_\omega\}$ and $Q_\omega(\omega') = P_{\omega,\omega'}$ for all ω'.

Definition 4.5. *A* log-linear Markov family *is a family of finite Markov chains X_t^β with $\beta \in [1, +\infty[$, defined on a common state space Ω, such that*

(i) the chain $X_t = X_t^1$ is irreducible and admits a decomposition $(M_\omega, q_\omega)_{\omega\in\Omega}$,

(ii) each X_t^β with $\beta > 1$ admits a decomposition $(M_\omega^\beta, q_\omega)_{\omega\in\Omega}$ given by

$$M_\omega^\beta = \{Q_\omega^\beta|Q_\omega \in M_\omega\}$$

where $\ln Q_\omega^\beta(\omega') = \beta\cdot\ln Q_\omega(\omega')$ whenever $Q_\omega(\omega') > 0$ (and $Q_\omega^\beta(\omega') = 0$ otherwise).

A log-linear Markov family can be seen as an interpolation between the X_t^1 chain (the "pure noise" chain) and a "limit chain" as $\beta \to \infty$. Irreducibility of the pure-noise chain implies irreducibility of all chains in the family, but not of the limit chain. A state ω is *stochastically stable* if $\lim_{\beta\to\infty} \mu^\beta(\omega) > 0$, where μ^β is the invariant distribution for $\beta > 0$.

Example 4.5. Consider the logit-response dynamics with revision process q. Its decomposition corresponds to equation 4.4. That is, the transition mechanisms Q_J available at a state s correspond to the revising sets J, and $Q_J(s') = e^{\beta \cdot U_J(s',s)}$. The pure-noise chain corresponds to the $\beta = 1$ case, and the limit chain is the best-response dynamics.

Given a log-linear family, a *transition tree* is defined analogously to a revision tree, i.e. a pair (T, γ) where T is a tree such that $(\omega, \omega') \in T$ only if $M_{\omega,\omega'} \neq \emptyset$ and $\gamma : T \mapsto M$ is such that $\gamma(\omega, \omega') \in M_{\omega,\omega'}$ for each $(\omega, \omega') \in T$. That is, γ selects a transition mechanism for each link in the tree. Denote the set of all transition ω-trees by $\mathcal{T}(\omega)$.

Analogously to Lemma 4.1 in the Appendix, straightforward but cumbersome computations allow to give an exact characterization of the invariant distribution $\mu^\beta(\omega)$. This in turn allows to establish the analogue of Theorem 4.2. The *waste* of a revision tree $(T, \gamma) \in \mathcal{T}(\omega)$ is defined as

$$W(T, \gamma) = \sum_{(\omega,\omega')\in T} \left(\max_{\omega''\in\Omega} Q_{\gamma(\omega,\omega')}(\omega'', \omega) \right) - Q_{\gamma(\omega,\omega')}(\omega', \omega).$$

The *stochastic potential* of a state ω is defined as $W(\omega) = \min_{(T,\gamma)\in\mathcal{T}(\omega)} W(T, \gamma)$.

Theorem 4.3. *Consider a log-linear Markov family. A state ω is stochastically stable if and only if it minimizes $W(\omega)$ among states.*

Log-linear Markov families can be used to analyze a large variety of learning models. Alternative transition mechanisms might not only correspond to alternative revision opportunities, as in the case of the logit-response dynamics, but can also be used to model different behavioral rules of the agents, for example imitation rules.

In this chapter, we have centered on the logit-response dynamics and hence it is natural to consider log-linear Markov families as a generalization. It would of course be possible to further generalize the framework to allow for perturbations which are not of the log-linear form. Such a framework would allow to encompass e.g. the mistakes model as a particular case (with the pure noise chain being the mutation process and the limit chain myopic best reply). Related approaches have been pursued by Myatt and Wallace (2003) and Beggs (2005), who consider families of Markov chains with transition probabilities P^β such that the limits $\lim_{\beta\to\infty} -\frac{1}{\beta} \ln P^\beta_{\omega,\omega'}$ are well-defined.[18]

4.6 Conclusions

The mistakes model of Kandori et al. (1993) and Young (1993) is analytically flexible due to the well-known graph-theoretic characterization of the stochastically stable states. It has been often criticized, e.g. by Bergin and Lipman (1996), due to the sensitivity of the results to the specification of the noise process. Other dynamics, like the logit-response dynamics of Blume (1993), present more solid foundations but analytical results can be derived only for particularly convenient frameworks.

Here we have presented a characterization of the stochastically stable states of a generalization of the logit-response dynamics. This new characterization is in the spirit of the mistakes model. We have illustrated the approach studying convergence to the set of Nash equilibria of the logit-response dynamics in general classes of games. Convergence obtains for best-response potential games but fails for generalized ordinal potential games. The selection of potential maximizers in exact potential games appears to be a fragile result, robust neither to generalizations of the considered game class nor to the specification of revision opportunities.

[18]Myatt and Wallace (2003) examine stochastic stability in a learning model where payoffs are perturbed by normally distributed shocks. They show that the addition of a strictly dominated strategy can change the selection result. Following the approach in Ellison (2000), Beggs (2005) uses graph-theoretic arguments to obtain general results on waiting times.

4.A Appendix

4.A.1 Proof of Theorem 4.1

The proof proceeds as follows. First, we introduce a few auxiliary concepts. Then we use these concepts to provide an exact (but cumbersome) characterization of the invariant distribution for fixed, finite β in Lemma 4.1. Last, we use this characterization to prove Theorem 4.1.

Given a graph G on the state space S, a mapping

$$\gamma : G \mapsto \mathcal{P}(I)$$

such that $\gamma(s, s') \in R_{s,s'}$ for each $(s, s') \in G$ is called a *revision selection for G*. For each transition in G, a revision selection for G picks exactly one of the possible revising sets making that transition (potentially) possible. Thus, a revision tree is a pair (T, γ) made of an s-tree involving only feasible transitions under the revision process q, and a revision selection γ for T.

Denote the set of all revision selections for a graph G by $\mathcal{S}(G)$.

Let $M = S \times \mathcal{R}^q$ denote the set of all pairs made of one state and one revising set. Consider a subset $N \subseteq M$. A *realization* r for N is a mapping $r : N \mapsto S$ such that $r(s, J) \in R_s^J$ for all $s \in S$ and all $J \in \mathcal{R}^q$, $J \neq \emptyset$. The set of all realizations for N is denoted $R(N)$. A *complete realization* is just a realization for M.

A *completion* of a revision tree (T, γ) is a complete realization such that $r(s, \gamma(s, s')) = s'$ for all $(s, s') \in T$. In words, a completion assigns a feasible outcome for each state and each possible revising set such that, whenever the revising set is the one specified by the selection for the (unique) arrow leaving the state in the tree, the outcome is precisely the state this arrow leads to. Let $\mathcal{C}(T, \gamma)$ be the set of all completions of (T, γ).

If γ is a revision selection for a tree T and $N^\gamma = \{(s, \gamma(s, s'))|(s, s') \in T\}$, then $R(N^\gamma)$ can be interpreted as the set of possible realizations of the selection γ.

Lemma 4.1 (The Decomposition Lemma)**.** *The invariant distribution μ of the logit-response dynamics with revision process q satisfies for each $s^* \in S$*

$$\mu(s^*) \propto \sum_{(T,\gamma)\in\mathcal{T}(s^*)} P(T,\gamma) \sum_{r\in\mathcal{C}(T,\gamma)} e^{Q(r)}, \tag{4.5}$$

where

$$P(T,\gamma) = \prod_{(s,s')\in T} q_{\gamma(s,s')}$$

and

$$Q(r) = \sum_{s\in S}\sum_{J\in\mathcal{R}^q} U_J(r(s,J),s).$$

Proof. Let $\mathcal{T}_0(s^*)$ be the set of all s^*-trees. By Freidlin and Wentzell (1988, Lemma 3.1),

$$\mu(s^*) \propto \sum_{T\in\mathcal{T}_0(s^*)} \prod_{(s,s')\in T} P_{s,s'}.$$

Note that s-trees including transitions which are not feasible under q contribute zero to the sum above. Fix a tree $T \in \mathcal{T}_0(s^*)$ such that all transitions are feasible under q. Using the decomposition of the transition probabilities (4.4),

$$\prod_{(s,s')\in T} P_{s,s'} = \prod_{(s,s')\in T}\left(\sum_{J\in R_{s,s'}} q_J \frac{e^{U_J(s',s)}}{\sum_{s''\in R_s^J} e^{U_J(s'',s)}}\right).$$

Expanding the RHS yields

$$\sum_{\gamma\in\mathcal{S}(T)}\left(\prod_{(s,s')\in T} q_{\gamma(s,s')} \frac{e^{U_{\gamma(s,s')}(s',s)}}{\sum_{s''\in R_s^{\gamma(s,s')}} e^{U_{\gamma(s,s')}(s'',s)}}\right) =$$

$$= \sum_{\gamma\in\mathcal{S}(T)}\left(P(T,\gamma)\frac{e^{\sum_{(s,s')\in T} U_{\gamma(s,s')}(s',s)}}{\prod_{(s,s')\in T}\sum_{s''\in R_s^{\gamma(s,s')}} e^{U_{\gamma(s,s')}(s'',s)}}\right).$$

Expanding the denominator in the last expression yields

$$\sum_{\gamma\in\mathcal{S}(T)}\left[P(T,\gamma)\, e^{\sum_{(s,s')\in T} U_{\gamma(s,s')}(s',s)}\left(\sum_{r\in R(N^\gamma)} e^{\sum_{(s,s')\in T} U_{\gamma(s,s')}(r(s,\gamma(s,s')),s)}\right)^{-1}\right].$$

We multiply and divide the last expression by

$$\sum_{r\in R(M\setminus N^\gamma)} e^{\sum_{(s,J)\in M\setminus N^\gamma} U_J(r(s,J),s)}$$

and obtain

$$\sum_{\gamma\in\mathcal{S}(T)}\left[P(T,\gamma)\left(\sum_{r\in\mathcal{C}(T,\gamma)}e^{\sum_{(s,J)\in M}U_J(r(s,J),s)}\right)\left(\sum_{r\in R(M)}e^{\sum_{(s,J)\in M}U_J(r(s,J),s)}\right)^{-1}\right].$$

The last term in brackets is a constant which is independent of both γ and T, and hence is irrelevant for proportionality of $\mu(s^*)$. The proof is completed observing that $\sum_{(s,J)\in M}U_J(r(s,J),s)=Q(r)$. □

We are now ready to prove Theorem 4.1, i.e. stochastically stable states are those where the waste is minimized across revision trees.

Proof of Theorem 4.1. Fix a revision process q. Let μ^β denote the invariant distribution of the logit response dynamics for noise level β. By Lemma 4.1 we have that, for every state s,

$$\mu^\beta(s)\propto\sum_{(T,\gamma)\in\mathcal{T}(s)}P(T,\gamma)\sum_{r\in\mathcal{C}(T,\gamma)}e^{\beta\cdot Q(r)},$$

where $Q(r)=\sum_{s\in S}\sum_{J\in\mathcal{R}^q}U_J(r(s,J),s)$, and $P(T,\gamma)>0$ for all (T,γ). As $\beta\to\infty$, only the completion r which maximizes $Q(r)$ among all completions for all revision trees $(T,\gamma)\in\mathcal{T}(s)$ matters for stochastic stability of state s, since its effect dominates for large β. Specifically, stochastically stable states, i.e. those satisfying that $\lim_{\beta\to\infty}\mu^\beta(s)>0$, are exactly those states $s\in S$ for which the expression

$$\max_{(T,\gamma)\in\mathcal{T}(s)}\max_{r\in\mathcal{C}(T,\gamma)}Q(r)$$

is maximal among all states.

For any given revision tree (T,γ), the completion r^* which maximizes $Q(r)$ among all completions $r\in\mathcal{C}(T,\gamma)$ clearly involves only best responses for all revising players on all pairs $(s,J)\notin N^\gamma$, i.e. in state-revising set pairs not used for transitions in the revision tree.

Let $r^{\max}$ be a complete realization involving only best responses. It follows that a state $s\in S$ maximizes $\max_{(T,\gamma)\in\mathcal{T}(s)}\max_{r\in\mathcal{C}(T,\gamma)}Q(r)$ if and only if it maximizes $\max_{(T,\gamma)\in\mathcal{T}(s)}\sum_{(s,s')\in T}\big(U_{\gamma(s,s')}(s',s)-U_{\gamma(s,s')}(r^{\max}(s,\gamma(s,s')),s)\big)$. Since $\sum_{(s,s')\in T}\big(U_{\gamma(s,s')}(s',s)-U_{\gamma(s,s')}(r^{\max}(s,\gamma(s,s')),s)\big)=-W(T,\gamma)$, it follows that stochastically stable states are those having minimal stochastic potential $\min_{(T,\gamma)\in\mathcal{T}(s)}W(T,\gamma)$. □

References of Chapter 4

ALÓS-FERRER, C. AND A. B. ANIA (2005): "The Evolutionary Stability of Perfectly Competitive Behavior," *Economic Theory*, 26, 497–516.

ALÓS-FERRER, C. AND G. KIRCHSTEIGER (2007): "Learning and Market Clearing," Mimeo.

BARON, R., J. DURIEU, H. HALLER, AND P. SOLAL (2002a): "A Note on Control Costs and Logit Rules for Strategic Games," *Journal of Evolutionary Economics*, 12, 563–575.

——— (2002b): "Control Costs and Potential Functions for Spatial Games," *International Journal of Game Theory*, 31, 541–561.

BEGGS, A. (2005): "Waiting Times and Equilibrium Selection," *Economic Theory*, 25, 599–628.

BERGIN, J. AND B. L. LIPMAN (1996): "Evolution with State-Dependent Mutations," *Econometrica*, 64, 943–956.

BLUME, L. (1993): "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, 5, 387–424.

——— (1997): "Population Games," in *The Economy as an Evolving Complex System II*, ed. by B. Arthur, S. Durlauf, and D. Lane, Addison-Wesley, Reading, MA, 425–460.

——— (2003): "How Noise Matters," *Games and Economic Behavior*, 44, 251–271.

ELLISON, G. (2000): "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution," *Review of Economic Studies*, 67, 17–45.

FREIDLIN, M. AND A. WENTZELL (1988): *Random Perturbations of Dynamical Systems, 2nd Ed.*, New York: Springer Verlag.

HART, S. AND A. MAS-COLELL (2000): "A Simple Adaptive Procedure Leading to Correlated Equilibrium," *Econometrica*, 68, 1127–1150.

HOFBAUER, J. AND W. SANDHOLM (2002): "On the Global Convergence of Stochastic Fictitious Play," *Econometrica*, 70, 2265–2294.

HOFBAUER, J. AND K. SIGMUND (1988): *The Theory of Evolution and Dynamical Systems*, Cambridge, UK: Cambridge University Press.

HOFBAUER, J. AND G. SORGER (1999): "Perfect Foresight and Equilibrium Selection in Symmetric Potential Games," *Journal of Economic Theory*, 85, 1–23.

KANDORI, M., G. J. MAILATH, AND R. ROB (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61, 29–56.

KANDORI, M. AND R. ROB (1995): "Evolution of Equilibria in the Long Run: A General Theory and Applications," *Journal of Economic Theory*, 65, 383–414.

MARUTA, T. (2002): "Binary Games with State Dependent Stochastic Choice," *Journal of Economic Theory*, 103, 351–376.

MATTSSON, L.-G. AND J. W. WEIBULL (2002): "Probabilistic Choice and Procedurally Bounded Rationality," *Games and Economic Behavior*, 41, 61–78.

MCKELVEY, R. D. AND T. R. PALFREY (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6–38.

MONDERER, D. AND L. SHAPLEY (1996): "Potential Games," *Games and Economic Behavior*, 14, 124–143.

MYATT, D. P. AND C. WALLACE (2003): "A Multinomial Probit Model of Stochastic Evolution," *Journal of Economic Theory*, 113, 286–301.

SANDHOLM, W. H. (1998): "Simple and Clever Decision Rules for a Model of Evolution," *Economics Letters*, 61, 165–170.

——— (2007): "Pigouvian Pricing and Stochastic Evolutionary Implementation," *Journal of Economic Theory*, 132, 367–382.

THURSTONE, L. (1927): "A Law of Comparative Judgement," *Psychological Review*, 34, 273–286.

VAN DAMME, E. AND J. W. WEIBULL (2002): "Evolution in Games with Endogenous Mistake Probabilities," *Journal of Economic Theory*, 106, 296–315.

VEGA-REDONDO, F. (1997): "The Evolution of Walrasian Behavior," *Econometrica*, 65, 375–384.

VOORNEVELD, M. (2000): "Best-Response Potential Games," *Economics Letters*, 66, 289–295.

YOUNG, P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.

——— (1998): *Individual Strategy and Social Structure*, Princeton, New Jersey: Princeton University Press.

Complete References

ABEL, A. (1988): "The Implications of Insurance for the Efficacy of Fiscal Policy," *Journal of Risk and Insurance*, 55, 339–378.

ACEMOGLU, D. AND J. ROBINSON (2000): "Why Did the West Extend the Franchise? Democracy, Inequality and Growth in Historical Perspective," *Quarterly Journal of Economics*, 115, 1167–1199.

AINSLIE, G. AND N. HASLAM (1992): "Hyperbolic Discounting," in *Choice Over Time*, ed. by G. Loewenstein and J. Elster, New York: Russell Sage Foundation.

ALÓS-FERRER, C. AND A. B. ANIA (2005): "The Evolutionary Stability of Perfectly Competitive Behavior," *Economic Theory*, 26, 497–516.

ALÓS-FERRER, C. AND G. KIRCHSTEIGER (2007): "Learning and Market Clearing," Mimeo.

BARON, R., J. DURIEU, H. HALLER, AND P. SOLAL (2002a): "A Note on Control Costs and Logit Rules for Strategic Games," *Journal of Evolutionary Economics*, 12, 563–575.

——— (2002b): "Control Costs and Potential Functions for Spatial Games," *International Journal of Game Theory*, 31, 541–561.

BECHARA, A., H. DAMASIO, D. TRANEL, AND A. DAMASIO (1997): "Deciding Advantageously Before Knowing the Advantageous Strategy," *Science*, 275, 1293–1295.

BECKER, G. (1968): "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 76, 169–217.

BEGGS, A. (2005): "Waiting Times and Equilibrium Selection," *Economic Theory*, 25, 599–628.

BENOÎT, J.-P. AND M. OSBORNE (1995): "Crime, Punishment and Social Expenditure," *Journal of Institutional and Theoretical Economics*, 151, 326–347.

BERGIN, J. AND B. L. LIPMAN (1996): "Evolution with State-Dependent Mutations," *Econometrica*, 64, 943–956.

BERGSTROM, T. (1997): "Storage for Good Times and Bad: Of Rats and Men," Mimeo.

BLUME, L. (1993): "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, 5, 387–424.

——— (1997): "Population Games," in *The Economy as an Evolving Complex System II*, ed. by B. Arthur, S. Durlauf, and D. Lane, Addison-Wesley, Reading, MA, 425–460.

——— (2003): "How Noise Matters," *Games and Economic Behavior*, 44, 251–271.

BOADWAY, R., M. LEITE-MONTEIRO, M. MARCHAND, AND P. PESTIEAU (2006): "Social Insurance and Redistribution with Moral Hazard and Adverse Selection," *Scandinavian Journal of Economics*, 108, 279–298.

BRENNAN, G. (1973): "Pareto Desirable Redistribution: The Non-Altruistic Dimension," *Public Choice*, 14, 43–67.

BUCHANAN, J. (1975): "The Political Economy of Franchise in the Welfare State," in *Capitalism and Freedom: Problems and Prospects*, ed. by R. Selden, Charlottesville: University Press of Virginia, 52–91.

COOPER, S. AND R. KAPLAN (1982): "Adaptive "Coin-flipping": A Decision-Theoretic Examination of Natural Selection for Random Individual Variation," *Journal of Theoretical Biology*, 94, 135–151.

CREMER, H. AND P. PESTIEAU (1996): "Redistributive Taxation and Social Insurance," *International Tax and Public Finance*, 3, 281–295.

CROCKER, K. AND A. SNOW (1985): "The Efficiency of Competitive Equilibria in Insurance Markets with Asymmetric Information," *Journal of Public Economics*, 26, 207–219.

CURRY, P. (2001): "Decision Making Under Uncertainty and the Evolution of Interdependent Preferences," *Journal of Economic Theory*, 98, 357–369.

CUTLER, D. AND J. GRUBER (1996): "Does Public Insurance Crowd Out Private Insurance?" *Quarterly Journal of Economics*, 111, 391–430.

DAMASIO, A. (1994): *Descartes' Error: Emotion, Reason, and the Human Brain*, New York: Avon.

DASGUPTA, P. AND E. MASKIN (2005): "Uncertainty and Hyperbolic Discounting," *American Economic Review*, 95, 1290–1299.

DEKEL, E., J. C. ELY, AND O. YILANKAYA (2007): "Evolution of Preferences," *Review of Economic Studies*, 74, 685–704.

DEMOUGIN, D. AND R. SCHWAGER (2000): "Excess Burden of Criminality and Redistribution," *International Review of Law and Economics*, 20, 329–351.

——— (2003): "Law Enforcement and Criminality: Europe vs. USA," *International Review of Law and Economics*, 23, 217–225.

EATON, J. AND H. S. ROSEN (1979): "Notes on Optimal Wage Taxation and Uncertainty," *NBER Working Paper No. 388*.

——— (1980): "Labor Supply, Uncertainty and Efficient Taxation," *Journal of Public Economics*, 14, 365–374.

ECKSTEIN, Z., M. EICHENBAUM, AND D. PELED (1985): "Uncertain Lifetimes and the Welfare Enhancing Properties of Annuity Markets and Social Security," *Journal of Public Economics*, 26, 303–326.

ELLISON, G. (2000): "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution," *Review of Economic Studies*, 67, 17–45.

ELSTER, J. (1985): "Weakness of Will and the Free-Rider Problem," *Economics and Philosophy*, 1, 231–265.

FLODEN, M. (2006): "Labor Supply and Saving under Uncertainty," *Economic Journal*, 116, 721–737.

FREDERICK, S. AND G. LOEWENSTEIN (1999): "Hedonic Adaption," in *Well-Being: The Foundations of Hedonic Psychology*, ed. by D. E. Kahneman, D. and N. Schwarz, New York: Russell Sage Foundation, 302–329.

FREDERICK, S., G. LOEWENSTEIN, AND T. O'DONOGHUE (2002): "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40, 351–401.

FREIDLIN, M. AND A. WENTZELL (1988): *Random Perturbations of Dynamical Systems, 2nd Ed.*, New York: Springer Verlag.

GAZZANIGA, M., R. IVRY, AND G. MANGUN (1998): *Cognitive Neuroscience: The Biology of the Mind*, New York: Norton.

GROSSMAN, H. (1995): "Robin Hood and the Redistribution of Income," *European Journal of Political Economy*, 11, 399–410.

——— (1998): "Producers and Predators," *Pacific Economic Review*, 3, 169–187.

——— (2001): "The Creation of Effective Property Rights," *American Economic Review, Papers and Proceedings*, 91, 347–352.

GROSSMAN, H. AND M. KIM (2003): "Educational Policy: Egalitarian or Elitist?" *Economics & Politics*, 15, 225–246.

GUL, F. AND W. PESENDORFER (2001): "Temptation and Self-control," *Econometrica*, 69, 1403–1435.

HANSSON, I. AND C. STUART (1999): "Malthusian Selection of Preferences," *American Economic Review*, 80, 529–544.

HART, S. AND A. MAS-COLELL (2000): "A Simple Adaptive Procedure Leading to Correlated Equilibrium," *Econometrica*, 68, 1127–1150.

HARTWICK, J. M. (2000): "Labor Supply under Wage Uncertainty," *Economics Letters*, 68, 319–325.

HECKMAN, J. (1993): "What Has Been Learned About Labor Supply in the Past Twenty Years?" *American Economic Review*, 83, 116–121.

HENRIET, D. AND J.-C. ROCHET (2004): "Is Public Health Insurance an Appropriate Instrument for Redistribution?" *Discussion Paper, GREMAQ, University of Toulouse*.

HOFBAUER, J. AND W. SANDHOLM (2002): "On the Global Convergence of Stochastic Fictitious Play," *Econometrica*, 70, 2265–2294.

HOFBAUER, J. AND K. SIGMUND (1988): *The Theory of Evolution and Dynamical Systems*, Cambridge, UK: Cambridge University Press.

HOFBAUER, J. AND G. SORGER (1999): "Perfect Foresight and Equilibrium Selection in Symmetric Potential Games," *Journal of Economic Theory*, 85, 1–23.

IMROHOROGLU, A., A. MERLO, AND P. RUPERT (2000): "On the Political Economy of Income Redistribution and Crime," *International Economic Review*, 41, 1–25.

KAHNEMAN, D. AND A. TVERSKY (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–291.

KANDORI, M., G. J. MAILATH, AND R. ROB (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61, 29–56.

KANDORI, M. AND R. ROB (1995): "Evolution of Equilibria in the Long Run: A General Theory and Applications," *Journal of Economic Theory*, 65, 383–414.

KIMBALL, M. S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica*, 58, 53–73.

KOÇKESEN, L., E. A. OK, AND R. SETHI (2000a): "Evolution of Interdependent Preferences in Aggregative Games," *Games and Economic Behavior*, 31, 303–310.

——— (2000b): "The Strategic Advantage of Negatively Interdependent Preferences," *Journal of Economic Theory*, 92, 274–299.

KOLMAR, M. AND D. BÖS (2003): "Aggressiveness and Redistribution," *Bonn Econ Discussion Papers*, 2/2003.

KUPFERMANN, I., E. KANDEL, AND S. IVERSON (2000): "Motivational and Addictive States," in *Principles of Neural Science*, ed. by E. Kandel, J. Schwartz, and T. Jessell, New York: McGraw-Hill, 998–1013, 4 ed.

LAIBSON, D. (1999): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–477.

LOW, H. AND D. MALDOOM (2004): "Optimal Taxation, Prudence and Risk-sharing," *Journal of Public Economics*, 88, 443–464.

MARUTA, T. (2002): "Binary Games with State Dependent Stochastic Choice," *Journal of Economic Theory*, 103, 351–376.

MATTSSON, L.-G. AND J. W. WEIBULL (2002): "Probabilistic Choice and Procedurally Bounded Rationality," *Games and Economic Behavior*, 41, 61–78.

MCCLURE, S., D. LAIBSON, G. LOEWENSTEIN, AND J. COHEN (2004): "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, 306, 503–507.

MCKELVEY, R. D. AND T. R. PALFREY (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6–38.

MIRRLEES, J. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38, 175–208.

MIYAZAKI, H. (1977): "The Rat Race and Internal Labor Markets," *Bell Journal of Economics*, 8, 394–418.

MONDERER, D. AND L. SHAPLEY (1996): "Potential Games," *Games and Economic Behavior*, 14, 124–143.

MULLIGAN, C. AND X. SALA-I-MARTIN (1999): "Social Security in Theory and Practice (II): Efficiency Theories, Narrative Theories, and Implications for Reform," *NBER Working Paper Series*, 7119.

MYATT, D. P. AND C. WALLACE (2003): "A Multinomial Probit Model of Stochastic Evolution," *Journal of Economic Theory*, 113, 286–301.

PARKER, S. C., Y. BELGHITAR, AND T. BARMBY (2005): "Wage Uncertainty and the Labor Supply of Self-employed Workers," *Economic Journal*, 115, 190–207.

POLLACK, R. (1968): "Consistent Planning," *Review of Economic Studies*, 35, 201–208.

PRATT, J. W. (1964): "Risk Aversion in the Small and in the Large," *Econometrica*, 32, 122–136.

RAYO, L. AND G. BECKER (2007a): "Evolutionary Efficiency and Happiness," *Journal of Political Economy*, 115, 302–337.

——— (2007b): "Habits, Peers, and Happiness: An Evolutionary Perspective," *American Economic Review, Papers and Proceedings*, 97, 487–491.

READ, D. (2001): "Intrapersonal Dilemmas," *Human Relations*, 54, 1093–1117.

READ, D., G. LOEWENSTEIN, AND M. RABIN (1999): "Choice Bracketing," *Journal of Risk and Uncertainty*, 19, 171–197.

RILEY, J. (1979): "Informational Equilibrium," *Econometrica*, 47, 331–359.

ROBSON, A. (1996): "A Biological Basis for Expected and Non-Excpected Utility," *Journal of Economic Theory*, 68, 397–424.

——— (2001a): "The Biological Basis of Economic Behavior," *Journal of Economic Literature*, 39, 11–33.

——— (2001b): "Why Would Nature Give Individuals Utility Functions?" *Journal of Political Economy*, 109, 900–914.

ROBSON, A. AND H. KAPLAN (2007): "Why We Grow Large and then Grow Old: Economics, Biology and Mortality," Mimeo.

ROBSON, A. AND L. SAMUELSON (2007): "The Evolution of Intertemporal Preferences," *American Economic Review*, 97, 496–500.

ROCHET, J.-C. (1991): "Incentives, Redistribution and Social Insurance," *Geneva Papers of Risk and Insurance*, 16, 143–165.

ROSS, S. A. (1981): "Some Stronger Measures of Risk Aversion in the Small and in the Large with Applications," *Econometrica*, 49, 621–638.

ROTHSCHILD, M. AND J. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay in the Economics of Incomplete Information," *Quarterly Journal of Economics*, 90, 629–649.

SAMUELSON, L. (2004): "Information-Based Relative Consumption Effects," *Econometrica*, 72, 93–118.

SANDHOLM, W. H. (1998): "Simple and Clever Decision Rules for a Model of Evolution," *Economics Letters*, 61, 165–170.

——— (2007): "Pigouvian Pricing and Stochastic Evolutionary Implementation," *Journal of Economic Theory*, 132, 367–382.

SANDMO, A. (1970): "The Effect of Uncertainty on Saving Decisions," *Review of Economic Studies*, 37, 353–360.

SCHELLING, T. (1984): "Self-Command in Practice, in Policy, and in a Theory of Rational Choice," *American Economic Review*, 74, 1–11.

SHEFRIN, H. AND R. THALER (1988): "The Behavioral Life-Cycle Hypothesis," *Economic Enquiry*, 26, 609–643.

SOZOU, P. (1998): "On Hyperbolic Discounting and Uncertain Hazard Rates," *Proceedings of the Royal Society of London Series B*, 265, 2015–2020.

SPENCE, A. (1978): "Product Differentiation and Performance in Insurance Markets," *Journal of Public Economics*, 10, 427–447.

STIGLITZ, J. E. (1982): "Utilitarianism and Horizontal Equity," *Journal of Public Economics*, 18, 1–33.

STROTZ, R. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165–180.

THALER, R. (1999): "Mental Accounting Matters," *Journal of Behavioral Decision Making*, 12, 183–206.

THURSTONE, L. (1927): "A Law of Comparative Judgement," *Psychological Review*, 34, 273–286.

TULLOCK, G. (1980): "Efficient Rent Seeking," in *Toward a Theory of the Rent-Seeking Society*, ed. by J. Buchanan, R. Tollison, and G. Tullock, College Station: Texas A&M University Press, 97–112.

van Damme, E. and J. W. Weibull (2002): "Evolution in Games with Endogenous Mistake Probabilities," *Journal of Economic Theory*, 106, 296–315.

Vega-Redondo, F. (1997): "The Evolution of Walrasian Behavior," *Econometrica*, 65, 375–384.

von Hayek, F. (1964): *The Constitution of Liberty*, Chicago: University of Chicago Press.

Voorneveld, M. (2000): "Best-Response Potential Games," *Economics Letters*, 66, 289–295.

Wilson, C. (1977): "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 12, 167–207.

Winston, G. (1980): "Addiction and Backsliding: A Theory of Compulsive Consumption," *Journal of Economic Behavior and Organization*, 1, 295–324.

Wright, S. (1932): "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution," *Proceedings of the VI International Congress of Genetics*, 1, 356–366.

Young, P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.

——— (1998): *Individual Strategy and Social Structure*, Princeton, New Jersey: Princeton University Press.

Erklärung

Ich versichere hiermit, dass ich die vorliegende Arbeit mit dem Thema

Risk, Behavior and Evolution

ohne unzulässige Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe. Die aus anderen Quellen direkt oder indirekt übernommenen Daten und Konzepte sind unter Angabe der Quelle gekennzeichnet. Weitere Personen, insbesondere Promotionsberater, waren an der inhaltlich materiellen Erstellung dieser Arbeit nicht beteiligt.[19] Die Arbeit wurde bisher weder im In- noch im Ausland in gleicher oder ähnlicher Form einer anderen Prüfungsbehörde vorgelegt.

Konstanz, den 24. Oktober 2007

(Nick Netzer)

[19] Siehe hierzu die Abgrenzung auf der folgenden Seite.

Abgrenzung

Ich versichere hiermit, dass ich Kapitel 2 und 3 der vorliegenden Arbeit ohne Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe.

Kapitel 1 entstammt einer gemeinsamen Arbeit mit Herrn Florian Scheuer (Massachusetts Institute of Technology). Die individuelle Leistung im Rahmen dieser Arbeit gliedert sich wie folgt:

i. Introduction/Conclusions: 50% Netzer / 50% Scheuer
ii. Labor Supply under Uncertainty: 60% Netzer / 40% Scheuer
iii. Optimal Taxation and Social Insurance: 50% Netzer / 50% Scheuer
iv. Imperfect Insurance Markets: 40% Netzer / 60% Scheuer

Kapitel 4 entstammt einer gemeinsamen Arbeit mit Herrn Carlos Alós-Ferrer (Universität Konstanz). Die individuelle Leistung im Rahmen dieser Arbeit gliedert sich wie folgt:

i. Introduction/Conclusions: 20% Netzer / 80% Alós-Ferrer
ii. The Logit-Response Dynamics: 50% Netzer / 50% Alós-Ferrer
iii. Stochastic Stability: 50% Netzer / 50% Alós-Ferrer
iv. Learning in Best-Response Potential Games: 60% Netzer / 40% Alós-Ferrer
v. Generalizations and Extensions: 40% Netzer / 60% Alós-Ferrer

Konstanz, den 24. Oktober 2007

(Nick Netzer)