

Efficiency and Equilibria in Games of Optimal Derivative Design*

Ulrich Horst

Santiago Moreno–Bromberg

Institut für Mathematik

Institut für Mathematik

Humboldt-Universität zu Berlin

Humboldt-Universität zu Berlin

Unter den Linden 6

Unter den Linden 6

10099 Berlin

10099 Berlin

horst@math.hu-berlin.de

smoreno@math.hu-berlin.de

May 10, 2010

Abstract

In this paper the problem of optimal derivative design, profit maximization and risk minimization under adverse selection when multiple agencies compete for the business of a continuum of heterogenous agents is studied. In contrast with the principal–agent models that are extended within, here the presence of ties in the agents’ best–response correspondences yields discontinuous payoff functions for the agencies. These discontinuities are dealt with via efficient tie–breaking rules. The main results of this paper are a proof of existence of (mixed–strategies) Nash equilibria in the case of profit–maximizing agencies, and of socially efficient allocations when the firms are risk minimizers. It is also shown that in the particular case of the entropic risk measure, there exists an efficient “fix–mix” tie–breaking rule, in which case firms share the whole market over given proportions.

PRELIMINARY - COMMENTS WELCOME

JEL classification: C62, C72, D43, D82, G14.

Keywords: Adverse selection, Nash equilibria, Pareto optimality, risk transfer, socially efficient allocations, tie–breaking rules.

*We thank Guillaume Carlier, Ivar Ekeland, and seminar participants at various institutions for valuable comments and suggestions. We are grateful to Alexander Fromm for his help in developing Section 5 of this work. This paper was finalized while the authors were visiting the Institute for Mathematical Sciences at the National University of Singapore. Financial support from the Deutsche Forschungsgemeinschaft through the SFB 649 “Economic Risk” and from the Alexander von Humboldt Foundation via a research fellowship is gratefully acknowledged.

1 Introduction

This work lies in the intersection of two fields that have produced a significant number of contributions in the last decade: Principal–Agent models under adverse selection and optimal risk sharing and the optimal design of financial products under translation invariant preferences.

In a Principal–Agent model of non–linear pricing of hedonic goods, a monopolist (the principal) has the capacity to deliver quality–differentiated products, which are assumed to lie on some compact and convex set $C \subset \mathbb{R}^n$. Such products are the *technologically feasible* goods from the point of view of the principal. These vectors are usually called *consumption bundles*, and each of its coordinates indicates how much a certain product possesses of a given attribute. In other words, goods are assumed to be fully described by a list of qualities that are relevant to the consumers. The buyers (or agents) whom the principal engages with the intent to trade have heterogenous preferences. This is captured by indexing the different characteristics or agent *types* with vectors in some set $\Theta \subset \mathbb{R}^m$, and including the types in the arguments of the buyers’ utility functions. The types are *private information*, i.e. the principal is aware of their statistical distribution μ , but she cannot distinguish an agent’s type prior to engaging him. The ill–informed principal designs *incentive compatible* catalogues with the intention of (at least partially) screening the market. The trading is done on a take–it–or–leave–it basis and it is assumed there is no second–hand market. The characterization of the solutions to problems of this kind can be found, among others, in the works of Armstrong [2], Mussa & Rosen [23] and Rochet & Choné [33].

The standard Principal–Agent model was extended by Carlier, Ekeland & Touzi in [9] to model a problem of optimal (over–the–counter) derivative design. They assume there is a direct cost to the principal when she delivers a derivative contract or financial product (a typical example are mutual funds) and that the agents’ utilities are of mean–variance type ¹. This allows them to phrase the principal’s profit maximization problem as a problem in the Calculus of Variations subject to convexity constraints, the latter capturing the incentive compatibility constraints on the set of admissible catalogues. This approach was further modified by Horst & Moreno–Bromberg in [20]. Here the authors assume the principal is exposed to some non–headgeable risky position, which she evaluates using a convex risk measure. This could be, for example, an insurer who has issued claims that are correlated to weather phenomena. The principal’s aim is to minimize her exposure by laying off part of her risk with heterogenous agents. To do so she proceeds as in [9] and designs a catalogue of derivatives written on her income, as well as a non–linear pricing schedule. In contrast with Carlier, Ekeland & Touzi, in this case the impact of each individual trade on the principal’s evaluation of her risk is non–linear, and an individual trade does not necessarily reduce the principal’s exposure. The main results of both papers are existence and characterization of direct revelation mechanisms that maximize the principal’s income ([9]) or minimize the principal’s risk assessment ([20]).

A more complex scenario, in which our current model is embedded, contemplates an oligopoly. The firms composing the latter have to take into account both the adverse selection problem and the competition among them. The first challenge that needs to be addressed when one studies existence of equilibria and/or Pareto optimal allocations is the fact that the Revelation Principle, an important simplifying ingredient in the Principal–Agent literature, may no longer be applied without loss of generality. In other words, some equilibrium allocations may not be implementable via direct revelation mechanisms. There are two ways to overcome this. The first one, introduced by Epstein & Peters [13], consists in enlarging the message space to include not only the agent types, but also a description of the market situations. Despite the theoretical value of this result, it is in general not practical for applications. Alternatively, Martimort &

¹These are type-dependent utility functions of the form $U(\theta, X) = G(X) + \theta \cdot F(X)$, where the asset–space is \mathcal{X} , $G : \mathcal{X} \rightarrow \mathbb{R}$ is linear and $F : \mathcal{X} \rightarrow \Theta$.

Stole [24] and Page Jr. & Monteiro [27] developed a multi-agency analogue of the Revelation Principle, the Delegation Principle, which allows enough simplification of the general non-linear pricing game to have a workable setting. A second challenge is the emergence of ties in the agents' best-response mappings. When agents are indifferent between contracting with different firms, then these firms' payoff functions may have discontinuities. This precludes the use of the classical results of Debreu, Glicksberg and Fan (see for example [18]) to guarantee the existence of Nash equilibria. Instead, much of the current literature on the existence of Nash equilibria in multi-firm-agent games relies on Reny's seminal paper [31], and the necessary conditions he presents for the existence of Nash equilibria in discontinuous games. The results of Bagh & Jofré in [4], and of Page Jr. & Monteiro in [28] and [29] provide testable conditions that allow for Reny's results to be used in multi-firm-agent games.

The second field that influenced this work could be broadly labeled "risk minimization and sharing" with coherent or, more generally, convex risk measures. The notion of *coherent risk measures* was axiomatized by Artzner, Delbaen, Ebert & Heath in [3] as an "acceptable" way to assess the riskiness of a financial position. It was then extended to *convex risk measures* by, among others, Föllmer & Schied [16] and Frittelli & Rosazza Gianin [17]. This theory has experienced an accelerated development, as evidenced by a large number of publications. Some quite natural questions to address, in a multitude of settings, are the existence and structure of risk-minimizing positions or of risk-sharing allocations that are either Pareto optimal or that constitute an equilibrium of some sort. The problem of optimal risk sharing for convex risk measures was first studied by Barrieu & El Karoui in [6]. They gave sufficient conditions for the risk-sharing problem with general state spaces to have a solution. Their conditions can be verified for the special case where the initial endowments are deterministic and the agents use modifications of the same risk measure; Jouini, Schachermayer & Touzi proved in [21] the existence of optimal risk-sharing allocations when the economic agents assess risk using convex risk measures which are *law-invariant*. The optimal allocations are Pareto optimal but not necessarily *individually rational*. That is, a cash transfer, also called the *rent of risk exchange*, could be necessary to guarantee that the outcome leaves all parties better off than they originally were. It should also be mentioned that the setting in [21] is over-the-counter (OTC) in nature, hence the question of efficiency and implementability cannot be left to the market. Implementability could depend on the presence of a social planner who enforces cash-transfer schemes or other policies that generate individually rational or socially optimal outcomes. In contrast, equilibrium models do not require the presence of a regulator to guarantee the implementability of efficient allocations. Implementability and efficiency is left up to the (financial) market. Equilibrium models of incomplete markets where agents use convex risk measures to evaluate their risk exposures were studied by, e.g., Filipović & Kupper in [14] in the static case and Cheridito et al. [8] in the dynamic one.

In this paper we study the problem of optimal derivative design, profit maximization and risk minimization under adverse selection when multiple agencies compete for the business of a continuum of heterogeneous agents. We first extend the model of profit maximization presented in [9] to a multi-firm one. Here is where the theory of games played under adverse selection comes into play. In order to use recent results from the theory of catalogue games, we assume that each firm's strategy set consists of the closed, convex hull of a finite number of basic products. We believe this is consistent with the idea that each firm is exposed to a direct cost when delivering financial products or derivative securities, since it must purchase the underlying assets that shall be structured into its' product line. In mathematical terms, such assumption renders the strategy sets compact with respect to the same topology for which the players' preferences are continuous (a necessary assumption), and therefore makes the game in hand compact. We show that when the game is played under a particular kind of tie-breaking rules, then it is uniformly payoff secure and reciprocally upper semicontinuous and hence has a mixed-strategies Nash equilibrium. Such rules, which paraphrasing [29] we call *efficient*, might not be implemented if not for the

presence of a social planner. The regulator also plays a key role in our multi–agency extension of the risk minimization model studied in [20]. In both cases the impact of a single trade on the firm’s risk evaluation is highly non–linear and depends strongly on the firm’s overall position.² This results in two considerable technical difficulties. First, there is no reason to expect (and in general it will not be the case) that the firms’ payoff functions will be uniformly payoff secure.³ Second, most of the current results ([4], [5], [18], [27], [25],...) on existence of Nash equilibria require some form of (quasi-) concavity (except [34] where a complete, yet very challenging to verify, characterization of Nash equilibria is presented). In our model these conditions are only satisfied by the mixed extension of the game. Even if the game were uniformly payoff secure and (weakly) reciprocally upper semicontinuous (by no means a given),⁴ considering a mixed extension of the game would mean the following: firms view linear aggregation of possible risk evaluations as the way to assess the influence of others’ in their own risk. Whether this approach is consistent with the ideas behind the theory of convex risk measure is in our opinion debatable. Taking the previous arguments into account, we do not seek the existence of (mixed–strategies) Nash equilibria. Our focus is instead on an existence proof for *socially efficient allocations* of risk exposures. Such risk allocations minimize the firms’ aggregate risk and can hence be thought of as the multi–firm–agent game analogous of the Pareto optimal allocations described in [21]. The proof of existence of efficient tie–breaking rules relies heavily on the fact that for fixed price schedules and tie–breaking rule, the contracts that minimize the aggregate risk can be expressed as the product of a type–independent random variable and a coefficient that depends exclusively on types. This separation result was already observed in [20], but it is in this paper where it is truly exploited. It is important to mention that efficient tie–breaking rules are generated endogenously. This implies that in our OTC model efficiency and regulation go hand–in–hand, even in the absence of a cash–transfer scheme. When it comes to implementability, we have obtained for the case of the entropic risk measure that among all efficient tie–breaking–rules, there is a constant ratio of market shares. In other words, there is a “fix–mix” ratio under which firms share the whole market, rather than segmenting it by agent types. The latter implies that all the firms offer the same indirect utility to each of the consumers. A real–world example where firms offer consumers essentially the same utility (using different products), and where the assumption of mean–variance optimizing consumers is appropriate, is retail banking. In such case “fix–mix” is also a reasonable assumption: retail banks differentiate customers according to their risk aversion when designing portfolio strategies, but do not necessarily try to appeal differently to customers with different attitudes towards risk. The methodology we used to obtain the “fix–mix” result suggests that it carries over to more general risk measures. A full analysis would be quite technical though, and certainly beyond the scope of this paper. We illustrate this result with a numerical example, in which we also find that (as expected) firms are worse off in the presence of competition, whereas the contrary is the case for the buyers. Moreover, the aggregate risk in the economy is better dealt with in the competitive, yet regulated, case. The numerical algorithm that we use for this example, a hybrid descent method, is also used to analyze an example where the firms are AV@R–minimizers. We find quite a sharp contrast in the firms’ risk profiles before and after trading when comparing the entropic and AV@R cases, which obeys the fact that when minimizing the latter, one should focus in the worst state of the world (our examples use, unavoidably, a finite probability space) in as much as the problem’s constraint allow, then move to the second worst and so on.

The remainder of this paper is organized as follows: Section 2 contains the description of our general framework. We leave the setting as general as possible while still describing the asymmetry of information

²In a model of profit maximization a firm contracts with some agent type *independently* of other types as long as that particular type contributes positively to the firm’s revenues.

³The “worst–case” tie–breaking rule, where each firm *assumes* ties will be broken in the most disadvantageous way for itself, does yield payoff security, but at the cost of reciprocally upper semicontinuity.

⁴This would imply the existence of Nash equilibria in mixed strategies

in the model, the best–response sets of the agents that give rise to ties, tie–breaking rules and the influence of the social planner. In Section 3 we study the game played among profit–maximizing firms, where the main result is the existence of mixed–strategies Nash equilibria. Section 4 is devoted to the risk–minimization game, where we prove the existence of socially efficient allocations. Finally, we present in Section 5 numerical algorithms to estimate equilibrium points and socially efficient allocations in some particular examples, as well as the “fix–mix” result mentioned above.

2 General Framework

We consider an economy that consists of two *firms* and a continuum of *agents*. We study both the case when the firms are profit maximizers, and the one where their objective is to minimize the risk assessments of some initial uncertain payoffs. The analysis of the two scenarios require different mathematical techniques and distinct notions of efficient allocations, so we study the two models separately.⁵

2.1 The financially feasible sets

The firms compete for the agents’ business by offering derivatives contracts. The set from which firm $i = 1, 2$ may choose products, i.e. its *financially feasible* set, is

$$\mathcal{X}_i \subset \mathbb{L}^2(\Omega, \mathcal{F}, \mathbb{P}),$$

where $(\Omega, \mathcal{F}, \mathbb{P})$ is a standard probability space. We assume these sets are closed, convex and bounded, and that $0 \in \mathcal{X}_i$. The boundedness of \mathcal{X}_i implies there is $M > 0$ such that $\|X_i\|_2 \leq M$ for all $X_i \in \mathcal{X}_i$. Any additional requirements on the sets \mathcal{X}_i shall be introduced when necessary.⁶ Throughout this paper we use the notation

$$X_n \xrightarrow{\|\cdot\|_2} X, \quad X_n \xrightarrow{a.s.} X \quad \text{and} \quad X_n \xrightarrow{w.} X$$

to indicate that the sequence $\{X_n\}$ converges strongly, almost surely and weakly to X . We also use $\mathbb{1}_A$ to denote the indicator function of a set A .

In the literature on multi–firm, non–linear pricing games, it is generally assumed that each firm chooses a compact subset \mathcal{Y}_i from its financially feasible set, and it devises a (non–linear) pricing schedule

$$p_i : \mathcal{Y}_i \rightarrow \mathbb{R}.$$

Each pair $(X, p_i(X))$ (where $X \in \mathcal{X}_i$) is called a *contract*. Then the Delegation and Competitive Taxation Principles (see for example [27]) allow for a without–loss–of–generality analysis of the existence of Nash equilibria by studying the (simpler) catalogue game played over compact subsets of the product–price space $\mathcal{X}_i \times \mathbb{R}$. In fact, the boundedness assumption on \mathcal{X}_i implies that prices belong to some compact set $P \subset \mathbb{R}$. We also study games played over catalogues, but their structure will be case–dependent. We write C_i to denote a catalogue offered by firm i , and (C_1, C_2) for a *catalogue profile* (also called a market situation). We postpone the specification of the criterions that the firms aim to optimize until Sections 3 and 4.

2.2 The agents’ preferences

The agents are heterogenous, mean–variance maximizers whose set of types (or characteristics) is

$$\Theta = [a, 1] \quad \text{for some} \quad a > 0.$$

⁵Our arguments can be extended to an economy with m firms. We chose to work on the case $m = 2$ for simplicity.

⁶We omit writing $i = 1, 2$ when we refer to properties shared by both firms.

The right endpoint of Θ has been normalized to 1, but it could be any finite value. What is required is for the set of types to be a compact subset of the strictly positive real numbers. An agent's type represents her risk aversion (hence the assumption $a > 0$ means there are no risk-neutral agents). In other words, given a contingent claim Y an agent of type θ assesses its worthiness via the (type-dependent) utility function

$$U(\theta, Y) = \mathbb{E}[Y] - \theta \text{Var}[Y].$$

The types are private information and not transparent to the firms. They are distributed according to a measure μ , which we assume is absolutely continuous with respect to Lebesgue measure. The measure μ is known to all firms. In other words, firms cannot distinguish an agent's type when engaging in trading with her (or they are legally prevented from doing so) but they know the overall distribution of types. This asymmetry of information, also known as *adverse selection*, prevents the firms from extracting all the above-reservation-utility wealth from each agent. The knowledge of μ is therefore essential for the agencies.⁷ In the upcoming sections we require the following auxiliary lemma:

Lemma 2.1 *The family of functions $\mathcal{U} = \{U(\theta, \cdot) \mid \theta \in \Theta\}$ is uniformly equicontinuous.*

PROOF. Let $\epsilon > 0$ and consider $X, X' \in \mathcal{X}$, and $\theta \in \Theta$, then

$$|U(\theta, X) - U(\theta, X')| = |\mathbb{E}[X] - \theta \text{Var}[X] - \mathbb{E}[X'] + \theta \text{Var}[X']|.$$

From the triangle inequality and the fact that $0 < \theta \leq 1$ we get

$$|U(\theta, X) - U(\theta, X')| \leq |\mathbb{E}[X] - \mathbb{E}[X']| + |\text{Var}[X] - \text{Var}[X']|.$$

It follows from the Cauchy-Schwartz inequality that

$$|\mathbb{E}[X] - \mathbb{E}[X']| \leq d_{\mathbb{P}}(X, X') \text{ and } |\text{Var}[X] - \text{Var}[X']| \leq d_{\mathbb{P}}(X, X') \left(\int_{\Omega} |X + X'| d\mathbb{P} + |\mathbb{E}[X] + \mathbb{E}[X']| \right),$$

where $d_{\mathbb{P}}(X, X')$ is the L^1 -distance with respect to \mathbb{P} . Since $-M \leq X, X' \leq M$, we obtain

$$|U(\theta, X) - U(\theta, X')| \leq d_{\mathbb{P}}(X, X') (1 + 4M) \leq (1 + 4M) \|X - X'\|_2.$$

Setting $\delta = \epsilon (1 + 4M)^{-1}$ yields the desired result. □

2.3 Indirect utilities and best-response sets

When an agent faces a catalogue profile (C_1, C_2) , she chooses a single contract from some C_i on a take-it-or-leave-it basis. Agents may choose any (available) contract they wish, but no bargaining regarding prices or products takes place. Given (C_1, C_2) , the indirect utility of an agent of type θ is:

$$v(\theta, C_1, C_2) \triangleq \sup \{U(\theta, X) - p \mid (X, p) \in C_i \text{ for some } i\}.$$

We assume that the agents have an outside option that yields their reservation utility, which we normalize to zero for all agents. For a fixed market situation, the function $v(\cdot, C_1, C_2) : \Theta \rightarrow \mathbb{R}$ is convex, since it is defined as the pointwise supremum of affine functions of θ . The presence of the outside option guarantees $v \geq 0$. The best-response set of the agents of type θ to a certain catalogue profile is defined as:

$$\mathcal{L}(\theta, C_1, C_2) \triangleq \text{argsup} \{U(\theta, X) - p \mid (X, p) \in C_i \text{ for some } i\}.$$

⁷We could make additional assumptions on μ and allow for $a = 0$.

In Sections 3 and 4 we make necessary assumptions on C_i as to guarantee $\mathcal{L}(\theta, C_1, C_2) \neq \emptyset$ for all $\theta \in \Theta$. Agents of type θ are indifferent among the elements of $\mathcal{L}(\theta, C_1, C_2)$, which they strictly prefer over any other contract that can be chosen from (C_1, C_2) . The Envelope Theorem (see for example [22]) implies that for any $(X(\theta), p(\theta)) \in \mathcal{L}(\theta, C_1, C_2)$

$$v'(\theta, C_1, C_2) = -\text{Var}[X(\theta)] \quad \mu - \text{a.s.} \quad (1)$$

In particular $v'(\cdot, C_1, C_2) \leq 0$ μ -a.s.. Equation (1) provides a valuable link between optimal contracts (from the point of view of the agents) and the indirect utility functions.

In order to exploit the information contained in μ , firms must choose catalogues that do not offer the agents incentives to lie about their types. If firm i intends to (partially) screen the market, the subset $C_i(\theta) \subset C_i$ of products from which it expects agents of type θ to make their choices must satisfy

$$E[X(\theta)] - \theta \text{Var}[X(\theta)] - p(\theta) \geq E[Y] - \theta \text{Var}[Y] - q, \quad \forall (X(\theta), p(\theta)) \in C_i(\theta), (Y, q) \in C_i.$$

Catalogues that satisfy this property are called *incentive compatible*. Let

$$\mathcal{L}(\theta, C_i) \triangleq \text{argsup}\{U(\theta, X) - p \mid (X, p) \in C_i\},$$

then C_i is incentive compatible if and only if $C_i(\theta) = \mathcal{L}(\theta, C_i)$.⁸ A catalogue where the products intended for agents of type θ yield at least their reservation utility is said to be *individually rational*. Since all reservation utilities are zero, an individually rational catalogue satisfies

$$E[X(\theta)] - \theta \text{Var}[X(\theta)] - p(\theta) \geq 0 \quad \forall (X(\theta), p(\theta)) \in \mathcal{L}(\theta, C_i).$$

In the presence of the outside option, participation in the market is endogenously determined.

2.4 Tie-breaking rules & market segmentation

A crucial element of multi-firm games is the presence of ties. There is no reason to assume that for a given catalogue profile (C_1, C_2) the sets $\mathcal{L}(\theta, C_1, C_2)$ ($\theta \in \Theta$) will be singletons. This fact renders the analysis considerably harder than it is for principal-agent games, where the Revelation Principle allows for such assumption. There can be both inter- and intra-firms ties, i.e. an agent may be indifferent between two products that are offered by distinct firms, or maybe between two contracts that are offered by the same agency. In what follows we study mechanisms via which ties are broken, as well as the partitions of Θ that are generated by tie-breaking. In order to study the ties that originate from a catalogue profile (C_1, C_2) , we define

$$v(\theta, C_i) = v_i(\theta) \triangleq \sup \{U(\theta, X) - p \mid (X, p) \in C_i\},$$

which are the one-catalogue analogues to $v(\theta, C_1, C_2)$. The v_i 's are related to $v(\theta, C_1, C_2)$ via

$$v(\theta, C_1, C_2) = \max \{v(\theta, C_1), v(\theta, C_2)\}.$$

Some important properties of the functions $v(\theta, C_i)$ are summarized in Lemma 2.2 below. In a nutshell it states that the indirect utility functions generated by incentive compatible catalogues are convex, and that there is a crucial link between the derivatives of such functions at a given type and the variance of the contracts that such type may choose.

⁸It has been shown in [33] that in general principal-agent models it is not possible to perfectly screen the market. In most instances there is a non-negligible set of agents who are pushed down to their reservation utilities (bunching of the first type), and another one where agents stay above their reservation levels, but they choose the same product despite the fact of their different preferences (bunching of the second type). In all likelihood, this behavior is inherited by multi-firm models.

Lemma 2.2 *Let (C_1, C_2) be a catalogue profile and assume $\mathcal{L}(\theta, C_i) \neq \emptyset$ for all $\theta \in \Theta$, then:*

1. *The functions $v(\cdot, C_i) : \Theta \rightarrow \mathbb{R}$ are convex.*
2. *If C_i is incentive compatible then $-\text{Var}[X_i(\theta)] \in \partial v(\theta, C_i)$ for all $(X_i(\theta), p_i(\theta)) \in C_i(\theta)$. In particular, the equation $-\text{Var}[X_i(\theta)] = v'(\theta, C_i)$ holds μ -a.s..*

PROOF. The mapping $\theta \mapsto U(\theta, X) - p$ is affine on the θ -th coordinate, so $v(\theta, C_i)$ is defined as the pointwise supremum of affine mappings and it is therefore convex (see for example Proposition 3.1 in [12]). By assumption $\mathcal{L}(\theta, C_i) \neq \emptyset$, an applying the Envelope Theorem as above we have that for any $(X_i(\theta), p_i(\theta)) \in \mathcal{L}(\theta, C_i)$, $-\text{Var}[X_i(\theta)] \in \partial v(\theta, C_i)$. The assumption of incentive compatibility implies that if $(X_i(\theta), p_i(\theta)) \in C_i(\theta)$, then $-\text{Var}[X_i(\theta)] = v'(\theta, C_i)$ μ -a.s. □

If for $\theta_0 \in \Theta$ we have $v_i(\theta_0) > v_{-i}(\theta_0)$, then the agents of type θ_0 will contract with firm i . The set of types that are indifferent between the firms' offers is

$$\Theta_0 \triangleq \{\theta \in \Theta \mid v(\theta, C_1) = v(\theta, C_2)\}.$$

To avoid ambiguities we assume that if $v_i(\theta) = v_{-i}(\theta) = 0$ then the corresponding agents opt for the outside option. The market is then segmented in the sets

$$\Theta_1 \triangleq \{v_1 > v_2\}, \quad \Theta_2 \triangleq \{v_2 > v_1\} \quad \text{and} \quad \Theta_0.$$

In order to deal with types in Θ_0 whose indirect utility is not zero, we define the set of tie-breaking rules as

$$F \triangleq \{f \in \mathbb{L}^0[a, 1] \mid 0 \leq f \leq 1\}.$$

From this point on, given a TBR f we write $f_1 = f$ and $f_2 = 1 - f$. Then $f_i \equiv 1$ on Θ_i , $f_i \equiv 0$ on Θ_{-i} , and for $\theta_0 \in \Theta_0$ the proportion $f_i(\theta_0)$ contracts with firm i .

In the sections below, we rephrase in as much as possible the interaction of the firms in terms of the indirect utilities generated by the catalogues they offer. This provides a clearer understanding of the market's segmentation, and in mathematical terms it allows us to use well established convex analysis machinery. To this end we define

$$\mathcal{C}_i \triangleq \left\{ v : \Theta \rightarrow \mathbb{R}_+ \mid v \geq 0, \exists \{X(\theta)\}_{\theta \in \Theta} \subset \mathcal{X}_i \text{ s.t. } v'(\theta) = -\text{Var}[X(\theta)] \right\}.$$

These are the sets of all possible (single-firm) indirect utilities that can be generated from incentive compatible catalogues contained in \mathcal{X}_i . The incentive compatibility is reflected in the requirement $v'(\theta) = -\text{Var}[X(\theta)]$, as in Lemma 2.2. We show below that these sets exhibit convenient compactness properties.

Proposition 2.3 *The sets \mathcal{C}_i are compact for the topology of uniform convergence.*

PROOF. When firm i designs a product line, it takes into account that

$$-M^2 \leq \text{E}[X] - \theta \text{Var}[X] \leq M.$$

This implies that all prices $p \in P$ must be below M to satisfy the individual rationality constraint. In counterpart $-M^2 \leq p$, otherwise all agents would be guaranteed an indirect utility above their reservation utility. Since $\|X\|_2 \leq M$ for all $X \in \mathcal{X}_i$, then $|v'_i| \leq M$ for all $v_i \in \mathcal{C}_i$, thus

$$0 \leq v_i \leq 3 \cdot \max\{M, M^2\}.$$

The convexity of the elements of the (closed) set \mathcal{C}_i implies they are locally Lipschitz (see for example [30]); moreover, the \mathbb{L}^2 -boundedness of \mathcal{X}_i together with Equation (1) imply the Lipschitz coefficients are uniformly bounded. This in turn means that \mathcal{C}_i is a bounded, closed and uniformly equicontinuous family, which by the Arzelà–Ascoli is then compact for the topology of uniform convergence. \square

2.5 The social planner

The fundamental theorems of welfare economics establish the equivalence between competitive equilibria (in complete markets) and efficiency, in the sense that frictionless competition leads to Pareto optimal allocations of resources and viceversa. In contrast, in OTC markets participants do not respond to given prices and therefore optimality is an inadequate notion to study. Instead, in the sections below we deal with the existence of *Nash equilibria* and *socially efficient allocations*.

Whereas in perfectly competitive settings market forces interact as to eventually reach efficient outcomes, OTC markets may require the influence of a *social planner* (a regulator) in order to achieve efficiency. This indirect market participant plays two important roles: First, he may choose to enforce certain kinds of TBRs in order to guarantee Nash and/or socially efficient outcomes. Second, he must make sure that individual rationality at the level of firms is preserved. This crucial implementability condition is a non-issue in competitive markets, where equilibrium allocations are also individually rational. Alternatively, given that socially efficient allocations are aggregately individually rational, the regulator could focus on efficiency and then establish a payment scheme among firms. These payment scheme would play the role that equilibrium prices do in competitive markets.

In what follows we assume that we work on regulated OTC markets, where the roles of the social planner are to seek that the market settles for socially (aggregate) optimal allocations, and that the latter are individually rational for the agencies.

3 Profit Maximization

In this section we analyze a non-cooperative game played among profit-maximizing firms under adverse selection. Our model is an extension into a multi-firm setting to the one studied in [9]. We show that if *efficient tie-breaking rules* are implemented (possibly through the influence of a regulator), then the game possesses a Nash equilibrium in mixed strategies.

3.1 The firms’ strategy sets & payoff functions

The theory of equilibria in multi-firm games developed (among others) in [24], [26] and [27] requires the strategy sets to be compact metric spaces. Moreover, the agents’ preferences must be continuous with respect to a topology that makes the strategy sets compact. In our view this implies that in general the strategy sets are actually closed and bounded sets of a finite-dimensional vector space. Following along the same lines, and with the aim of describing a profit-maximization model of OTC trading of derivatives contracts, we suppose that firms structure the latter from a set of “basic” products. These products are available in markets to which agents do not have access, and could be, for instance, OTC markets between “high-rollers”. The firms are then exposed to a cost when delivering each of these products.

We assume \mathcal{X}_i is the closed convex hull of a finite number of basic products $\{X_i^1, \dots, X_i^{m_i}\}$, and that the cost to firm i of delivering $X_i \in \mathcal{X}_i$ is given by a lower semicontinuous function

$$K_i : \mathcal{X}_i \rightarrow \mathbb{R}.$$

The per-contract profit of firm i when it sells claim X_i given the price schedule $p_i : \mathcal{X}_i \rightarrow \mathbb{R}$ is

$$p_i(X_i) - K_i(X_i).$$

Analyzing a multi-firm game played over non-linear price schedules would be daunting at best. Here the Revelation Principle (see for example [27]) comes to our rescue. In a nutshell, it states that the game played over product-price catalogues, i.e. elements of

$$\mathcal{P}_i \triangleq \{C_i \subset \mathcal{X}_i \times P \mid C_i \text{ closed}\},$$

is as general as the one played over closed subsets of \mathcal{X}_i and non-linear price schedules. We recall that $P \subset \mathbb{R}$ is the compact set where feasible prices lie. We assume that firms compete for the agents' business by offering product-price catalogues. In other words, the strategy sets \mathcal{P}_i are the sets of all compact subsets of $\mathcal{X}_i \times P$. We endow \mathcal{P}_i with the Hausdorff metric h . Since $\mathcal{X}_i \times P$ is a compact metric space, so is (\mathcal{P}_i, h) (see, for example [1], Section 3.17); furthermore, Tychonoff's theorem guarantees that $\mathcal{P} \triangleq \mathcal{P}_1 \times \mathcal{P}_2$ with the corresponding product metric h_p is also compact and metric. We write $\mathcal{B}(\mathcal{P})$ for the Borel σ -algebra in \mathcal{P} . The following proposition allows us to substitute, without loss of generality, a catalogue profile (C_1, C_2) for the agents' optimal choices associated with the market situation (C_1, C_2) , i.e., for

$$\left(\bigcup_{\theta} \mathcal{L}(\theta, C_1), \bigcup_{\theta} \mathcal{L}(\theta, C_2) \right)$$

as these sets are closed and hence compact.

Proposition 3.1 *For any catalogue $C_i \in \mathcal{P}_i$, the set $\bigcup_{\theta} \mathcal{L}(\theta, C_i)$ is closed.*

PROOF. Consider $(\bar{X}, \bar{p}) \in \text{cl}(\bigcup_{\theta} \mathcal{L}(\theta, C_i))$, then there exists a sequence $\{(X_n, p_n)\} \subset \bigcup_{\theta} \mathcal{L}(\theta, C_i)$ such that

$$(X_n, p_n) \xrightarrow{\|\cdot\|_2 \times |\cdot|} (\bar{X}, \bar{p}).$$

By construction $(X_n, p_n) \in \mathcal{L}(\theta, C_i)$ for some $\theta \in \Theta$, call it θ_n . Passing to a subsequence if necessary we may assume there is $\bar{\theta} \in \Theta$ such that $\theta_n \rightarrow \bar{\theta}$ and $(X_n, p_n) \rightarrow (\bar{X}, \bar{p})$ pointwise. If $(\bar{X}, \bar{p}) \in \bigcup_{\theta} \mathcal{L}(\theta, C_i)$ we are done, so let us assume the contrary, thus

$$\text{E}[\bar{X}] - \bar{\theta} \text{Var}[\bar{X}] - \bar{p} < \text{E}[X] - \bar{\theta} \text{Var}[X] - p \quad \forall (X, p) \in \mathcal{L}(\bar{\theta}, C_i). \quad (2)$$

By definition

$$\text{E}[X_n] - \theta_n \text{Var}[X_n] - p_n \geq \text{E}[X] - \theta_n \text{Var}[X] - p \quad \forall (X, p) \in C_i.$$

However, the (a.s.) convergence of the sequence $\{(X_n, p_n)\}$ implies

$$\text{E}[\bar{X}] - \bar{\theta} \text{Var}[\bar{X}] - \bar{p} \geq \text{E}[X] - \bar{\theta} \text{Var}[X] - p \quad \forall (X, p) \in C_i,$$

which would imply $(X, p) \in \mathcal{L}(\bar{\theta}, C_i)$, contradicting Equation (2). □

The maximal attainable profit for firm i from type θ given the (incentive compatible) catalogue profile (C_1, C_2) is

$$\pi_i(\theta, C_1, C_2) \triangleq \max \{ (p - K(X)) \mathbb{1}_{\{v_i \geq v_{-i}\}}(\theta) \mid (X, p) \in C_i \}.$$

Here the functions v_i are as defined in Section 2.4. The functions π_i are the building blocks of the firms' payoff functions and enjoy the following important continuity property:

Proposition 3.2 (Page Jr. [26]) *The maps $\pi_i : \Theta \times \mathcal{P} \rightarrow \mathbb{R}$ are upper semicontinuous on \mathcal{P} and $\mathcal{B}(\Theta) \times \mathcal{B}(\mathcal{P})$ -measurable.*

For a given TBR f , the payoff function of firm i is given by

$$\Pi_i(C_1, C_2) \triangleq \int_{\Theta} \pi_i(\theta, C_1, C_2) f_i(\theta) \mu(d\theta).$$

These payoff functions exhibit a highly discontinuous behavior, due to the presence of ties. However, the effect of each type-wise trade on the overall income of the firm is linear and easily quantifiable. Moreover, only agents who make a positive contribution to a firm's revenues contract with it. This last fact plays a role in the analysis of *payoff security* contained in Section 3.3.2.

3.2 Efficient tie-breaking rules & reciprocal upper semicontinuity

Roughly speaking, a TBR is said to be *efficient* if an agent contracts with a firm if and only if the latter values her as a customer at least as much as the competition. As we show below, this is equivalent to saying that an efficient TBR maximizes the aggregate profit for a given market situation (C_1, C_2) .

Definition 3.3 *Given a catalogue profile (C_1, C_2) , a tie-breaking rule $f \in F$ is called **efficient** if*

- $\pi_i(\theta, C_1, C_2) \geq \pi_{-i}(\theta, C_1, C_2)$ implies $f_i(\theta) > 0$,
- $\pi_i(\theta, C_1, C_2) < \pi_{-i}(\theta, C_1, C_2)$ implies $f_i(\theta) = 0$.

The definition above is equivalent to saying that f^* is efficient given the catalogue profile (C_1, C_2) if and only if

$$\sum_i \Pi_i(f^*, C_1, C_2) = \sup_{f \in F} \sum_i \Pi_i(f, C_1, C_2),$$

which shows that in general efficient TBRs are endogenously determined. In mathematical terms, efficient TBRs yield certain upper semicontinuity properties to the payoff functions of the firms, which are necessary to guarantee the existence of Nash equilibria.

Definition 3.4 *A game $\{(\Pi_i, \mathcal{P}_i)\}$ is said to be **reciprocal upper semicontinuous (RUSC)** for a given TBR f if the mapping*

$$(C_1, C_2) \mapsto \sum_i \Pi_i(f, C_1, C_2)$$

is upper semicontinuous.

The notion of RUSC games was introduced by Dasgupta and Maskin in [10] (labeled as complementary discontinuous or u.s.c-sum games) and later generalized by Reny in [31] in order to prove the existence of Nash equilibria in certain discontinuous games. Intuitively, in a RUSC game firm i can only approximate the payoff corresponding to a market situation (C_1, C_2) by actually playing C_i . We should note that Reny's definition is slightly stronger, but the one above (from [10]) is sufficient for our needs. Furthermore, Proposition 5.1 in [31] tells us that RUSC is inherited by the mixed extension of a game.

Lemma 3.5 (Page Jr. & Monteiro [28]) *If the game $\{(\mathcal{P}_i, \Pi_i)\}$ is played using efficient TBRs, then it is RUSC.*

PROOF. Consider the mapping $(C_1, C_2) \mapsto \sum_i \Pi_i(f^*, C_1, C_2)$. If f^* is efficient then we have

$$\begin{aligned}
\sum_i \Pi_i(f^*, C_1, C_2) &= \sup_{f \in F} \sum_i \Pi_i(f, C_1, C_2) \\
&= \sup_{f \in F} \sum_i \int_{\Theta} \pi_i(\theta, C_1, C_2) f_i(\theta) \mu(d\theta) \\
&= \int_{\Theta} \max_i \pi_i(\theta, C_1, C_2) \sum_i f_i(\theta) \mu(d\theta) \quad \text{for any } f \in F \\
&= \int_{\Theta} \max_i \pi_i(\theta, C_1, C_2) \mu(d\theta).
\end{aligned}$$

It follows from Proposition 3.2, that $(C_1, C_2) \mapsto \max_i \pi_i(\theta, C_1, C_2)$ is upper semicontinuous, hence so is $\sum_i \Pi_i(f^*, C_1, C_2)$. □

Notice that the definition of efficient TBR, together with the linear aggregation of type-wise profits, implies that the mapping $(C_1, C_2) \mapsto \sum_i \Pi_i(f, C_1, C_2)$ is independent of the TBR played if the latter it is efficient.

3.3 Existence of Nash equilibria

In this section we study necessary conditions for the existence of Nash equilibria in the agencies' game. We cannot employ the canonical results of Debreu, Glicksberg and Fan, given the discontinuities introduced by the TBRs. Instead we rely on the results of Page Jr. & Monteiro [27] and Reny [31]. Since the payoff functions Π_i are not quasiconcave, one cannot prove the existence of pure-strategies equilibria using the state-of-the-art results available in the literature. Instead we analyze the existence of mixed-strategies equilibria. In order to do so we show *uniform payoff security* of the game. This notion allows one to conveniently test the *payoff security* of a game's mixed extension.

3.3.1 Defining the mixed-strategies game

To define a *mixed catalogue game*, we let $\mathcal{M}(\mathcal{P}_j)$ be the set of probability measures on $(\mathcal{P}_j, \mathcal{B}(\mathcal{P}_j))$. We endow the *mixed catalogue strategy sets* $\mathcal{M}(\mathcal{P}_i)$ with the $\sigma(X^*, X)$ -topology generated by the dual pair $\langle \mathbb{L}^0(\mathcal{P}_i, \mathcal{B}(\mathcal{P}_i)), \mathcal{M}(\mathcal{P}_i) \rangle$. The game is played as follows: Each agency chooses an element $\lambda_i \in \mathcal{M}(\mathcal{P}_i)$, and its profit under the profile $\lambda = (\lambda_1, \lambda_2)$ is given by

$$\mathbf{\Pi}_i(\lambda) \triangleq \int_{\mathcal{P}} \Pi_i(C_1, C_2) \lambda(dC), \tag{3}$$

where $\lambda(dC) = \lambda_1(dC_1) \lambda_2(dC_2)$ is the corresponding product measure. By Proposition 3.2 the integrand is measurable, thus expression (3) is well defined. The game $\{(\mathcal{M}(\mathcal{P}_i), \mathbf{\Pi}_i)\}$ is the *mixed catalogue game* that extends $\{(\mathcal{P}_i, \Pi_i)\}$.

Definition 3.6 *A mixed profile λ is a Nash equilibrium for the game $\{(\mathcal{M}(\mathcal{P}_i), \mathbf{\Pi}_i)\}$ if*

$$\mathbf{\Pi}_i(\lambda_i, \lambda_{-i}) \geq \mathbf{\Pi}_i(\lambda'_i, \lambda_{-i}) \quad \text{for all } \lambda'_i \in \mathcal{M}(\mathcal{P}_i).$$

A Nash equilibrium in mixed strategies for $\{(\mathcal{M}(\mathcal{P}_j), \mathbf{\Pi}_i)\}$ can be viewed as an equilibrium in pure strategies for its mixed extension.

3.3.2 Uniform payoff security

A game in which the players' strategy sets are compact subsets of a topological space and their payoff functions are bounded is called *compact*. The mixed extension of the game in hand is a compact game. This is due to the fact that the sets $\mathcal{M}(\mathcal{P}_i)$ are closed, convex and bounded subsets of $(\mathbb{L}^0(\mathcal{P}_i, \mathcal{B}(\mathcal{P}_i)))^*$, thus by the Banach–Alaoglu theorem they are $\sigma(X^*, X)$ -compact. Reny provides sufficient conditions for existence of a Nash equilibrium in a compact game in the presence of discontinuous payoff functions. A key requirement is that the game is *payoff secure*. Intuitively this means that small deviations on the part of firm i 's competitors can be answered to with a deviation that leaves firm i within a small interval of its profit prior to deviations.

Definition 3.7 *Given the payoffs $\Pi_i : \mathcal{P}_i \rightarrow \mathbb{R}$, the catalogue game $\{(\mathcal{P}_i, \Pi_i)\}$ is **payoff secure** if for all $(C_1, C_2) \in \mathcal{P}$, and $\epsilon > 0$ there exist $C_i^* \in \mathcal{P}_i$ and $\delta > 0$ such that*

$$\Pi_i(C_i^*, C'_{-i}) \geq \Pi_i(C_i, C_{-i}) - \epsilon$$

for all $C'_{-i} \in \{D \in \mathcal{P}_{-i} \mid h_p(D, C_{-i}) < \delta\}$.

Payoff security of a game does not imply the same property for its mixed extension; however, such is the case with *uniform payoff security* (see Theorem 1 in [28]). Proving uniform payoff security of a game is simpler than dealing with the weak*-topology to show payoff security of its mixed extension.

Definition 3.8 *The catalogue game $\{(\mathcal{P}_i, \Pi_i)\}$ is **uniformly payoff secure** if for all $i = 1, 2$, $C_i \in \mathcal{P}_i$ and $\epsilon > 0$ there exist C_i^* such that for all $C_{-i} \in \mathcal{P}_{-i}$ there exists $\delta > 0$ that satisfies*

$$\Pi_i(C_i^*, C'_{-i}) \geq \Pi_i(C_i, C_{-i}) - \epsilon$$

for all $C'_{-i} \in \{D \in \mathcal{P}_{-i} \mid h_p(D, C_{-i}) < \delta\}$.

Proposition 3.9 *The game $\{(\mathcal{P}_i, \Pi_i)\}$ is uniformly payoff secure (hence the game $\{(\mathcal{M}(\mathcal{P}_j), \Pi_i)\}$ is payoff secure).*

PROOF. Let $C_i \in \mathcal{P}_i$ and $\epsilon > 0$ and define

$$C_i^\epsilon \triangleq \{(X, p - \epsilon) \mid (X, p) \in C_i \text{ and } p - \epsilon \geq K_i(X)\} \cup \{(0, 0)\},$$

i.e. C_i^ϵ is obtained from C_i by keeping all contracts that allow for an ϵ -decrease in prices without this rendering a negative price-cost benefit. All other contracts are replaced by $\{(0, 0)\}$. Lemma 2.1 implies there is $\delta > 0$ such that if $d((X, p), (X', p')) < \delta$ (where $d(\cdot, \cdot)$ is the distance generated by $\|\cdot\|_2 \times |\cdot|$) then for all $\theta \in \Theta$

$$|U(\theta, X) - p - U(\theta, X') + p'| < \epsilon. \quad (4)$$

Assume first that $\pi_i(\theta, C_i, C_{-i}) > \epsilon$ and let $(X, p) \in C_i$ be such that $U(\theta, X) - p = \pi_i(\theta, C_i, C_{-i})$. Then $(X, p - \epsilon) \in C_i^\epsilon$. Now consider C'_{-i} such that $h_p(C_{-i}, C'_{-i}) < \delta$, and let $(X', p') \in C'_{-i}$. By definition there exist $(Y, q) \in C_{-i}$ such that $d((Y, q), (X', p')) < \delta$, and by Equation (4)

$$U(\theta, X) - (p - \epsilon) \geq U(\theta, Y) - (q - \epsilon) > U(\theta, X') - p'.$$

The latter implies that $(X', p') \notin \mathcal{L}(\theta, C_i^\epsilon, C'_{-i})$ and in turn that an efficient TBR f for (C_i^ϵ, C'_{-i}) satisfies $f_i(\theta) = 1$. Thus

$$\pi_i(\theta, C_i^\epsilon, C'_{-i}) f_i(\theta) \geq \pi_i(\theta, C_i, C'_{-i}).$$

The inequality above is trivially fulfilled if $\pi_i(\theta, C_i, C_{-i}) \leq \epsilon$. We have that for any deviation C'_{-i} such that $h_p(C_{-i}, C'_{-i}) < \delta$ and for any efficient TBR f

$$\begin{aligned} \Pi_i(C_i^\epsilon, C'_{-i}) &\geq \int_{\Theta} \pi_i(\theta, C_i^\epsilon, C'_{-i}) f_i(\theta) \mu(d\theta) - \epsilon \\ &= \Pi_i(C_i, C_{-i}) - \epsilon. \end{aligned}$$

□

The proof of Proposition 3.9 is a particular case of the proof of Theorem 1 in [29]. We have included it here because it showcases how the linear aggregation of type-wise profits is key to the existence of Nash equilibria. The simple yet far reaching construction of C_i^ϵ is only possible in such case. Our result on existence of Nash equilibria relies on the following

Theorem 3.10 (Reny [31]) *Suppose that $G = \{(X_i, u_i)\}_{i=1}^m$ is a compact game which is also quasiconcave, reciprocally upper semicontinuous and payoff secure, then it possesses a pure-strategies Nash equilibrium.*

It follows from Lemma 3.5 that the (compact) game $\{(\mathcal{M}(\mathcal{P}_i), \Pi_i)\}$ is RUSC. The payoff functions are linear in each firm's strategy, hence quasiconcave. Proposition 3.9 then allows us to apply Theorem 3.10 to the mixed extension of the game $\{(\mathcal{P}_i, \Pi_i)\}$ to obtain the following

Theorem 3.11 *The game $\{(\mathcal{P}_i, \Pi_i)\}$ possesses a mixed-strategies Nash equilibrium.*

4 Risk Minimization

In this section we analyze the risk minimization problem faced by firms that have some initial risky endowments. The goal of the firms is to lay off parts of their risk on individual agents. This is done via OTC trading of derivatives contracts. The model presented here is the multi-firm version of the one introduced in [20]. In contrast with the previous section, we do not seek to prove the existence of Nash equilibria, but rather of *socially efficient allocations*. The main difficulty towards guaranteeing Nash equilibria stems from the non-linear, per-type impact on the firms' risk assessments. In the previous sections, the fact that any agent contracting with an agency had a positive impact on its revenues was essential to show uniform payoff security and RUSC. This is no longer the case for risk-minimizing firms. Here the risk associated to the aggregate of a firm's initial position plus a contract could be lower than the firm's initial risk assessment, but once all positions are taken into account, the firm might be better off not engaging the agents who would have chosen such contract.

4.1 The firms' strategy sets & risk assessments

The initial risky endowment of each firm is represented by a random variable $W_i \in \mathbb{L}^2(\Omega, \mathcal{F}, \mathbb{P})$. Firm i assesses its risk exposure using a convex and law invariant risk measure

$$\varrho_i : \mathbb{L}^2(\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \mathbb{R} \cup \{+\infty\},$$

which has the Fatou property. We refer the reader to Appendix C for a discussion on these maps, as well as related references. We restrict the firms' choices to catalogues of the form

$$C_i = \{(X_i(\theta), p_i(\theta))\}_{\theta \in \Theta}.$$

These are not direct revelation mechanisms: even if firm 1 were to offer the individually rational and incentive compatible catalogue C_1 , the presence of firm 2 might dissuade agents of type θ_0 from choosing $(X_1(\theta_0), p_1(\theta_0))$. For a given catalogue profile (C_1, C_2) and a given TBR f , the position of firm i after trading is

$$W_i - \int_{\Theta_i} (p_i(\theta) - X_i(\theta))\mu(d\theta) - \int_{\Theta_0} (p_i(\theta) - X_i(\theta))f_i(\theta)\mu(d\theta),$$

where Θ_i and Θ_0 are as defined in Section 2.4. If we write $p_i(\theta) = \theta v'_i(\theta, C_1, C_2) - v_i(\theta, C_1, C_2)$, and we rename $X(\theta) = X(\theta) - E[X(\theta)]$,⁹ then firm i 's risk assessment after trading is

$$A_i(C_1, C_2, f) \triangleq R_i(C_1, C_2, f) - I_i((C_1, C_2, f),$$

where

$$R_i(C_1, C_2, f) \triangleq \varrho_i \left(W_i - \int_{\Theta_i} X_i(\theta)\mu(d\theta) - \int_{\Theta_0} X_i(\theta)f_i(\theta)\mu(d\theta) \right),$$

denotes firm i 's risk and the associated income is given by

$$I_i((C_1, C_2, f) \triangleq \int_{\Theta_i} (\theta v'_i(\theta) - v_i(\theta))\mu(d\theta) + \int_{\Theta_0} (\theta v'_i(\theta) - v_i(\theta))f_i(\theta)\mu(d\theta).$$

As we have seen before, when a catalogue profile (C_1, C_2) is presented to the agents, the corresponding indirect utility functions (v_1, v_2) show how the market is segmented. The more interesting set is Θ_0 , i.e. the set of indifferent agents. Within this set, there are two intrinsically different situations. First, v_1 and v_2 may be identical over a non-negligible subset. These could be regarded as “true” ties, in the sense that they will have different impacts on the firms’ risk assessments under different TBRs. Second, there could be types for which v_1 and v_2 are secant. We show below that given our assumptions on μ , these types do not have a direct impact on the aggregate incomes (in fact they do not impact the A_i 's, but we do not require this for our arguments). However, they do indicate the points where agents switch from contracting with one firm to the other one, in other words they show where the customers’ preferences shift.

Definition 4.1 *Given a pricing schedule $(v_1, v_2) \in \mathcal{C}_1 \times \mathcal{C}_2$, we say that there is a **shift in the customers’ preferences** at $\theta \in \text{int}(\Theta)$ if*

$$(v_1(\theta_-) - v_2(\theta_-)) \cdot (v_1(\theta_+) - v_2(\theta_+)) < 0.$$

Proposition 4.2 *For any pricing schedule $(v_1, v_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ the set Θ_s of shifts in the customers’ preferences satisfies:*

1. Θ_s has at most countably many elements.
2. The derived set of Θ_s has measure zero.

PROOF. The set $\Theta_s(v_1, v_2)$ is the union of the closures of the sets

$$\Theta_s^1 \triangleq \{\theta \in \Theta \mid v_1(x) = v_2(x), v'_1(x) \neq v'_2(x)\}$$

and

$$\Theta_s^\epsilon \triangleq \{\theta \in \Theta \mid v_1(x) = v_2(x), v'_1(x) = v'_2(x), |v_1(\tilde{x}) - v_2(\tilde{x})| > 0 \forall \tilde{x} \in B(\epsilon, x) \setminus \{x\}, \epsilon > 0\}.$$

⁹This is possible due to the translation invariance of ϱ_i , and fully characterizes p_i in terms of v_i .

The set Θ_s^1 is denumerable and nowhere dense. This follows directly from the convexity of v_1 and v_2 and the fact that the corresponding supporting planes to $\text{graph}\{v_1\}$ and $\text{graph}\{v_2\}$ at $(\theta, v_1(\theta))$ are secant. By definition Θ_2 is denumerable and nowhere dense, since any two of its elements can be separated. Therefore $\text{cl}\{\Theta_s^1\}$ and $\text{cl}\{\Theta_s^\epsilon\}$ are themselves denumerable and nowhere dense. \square

As a result of the preceding proposition we see that the set of pre-images of the crossings of the graphs of two functions $v_1 \in \mathcal{C}_i$ and $v_2 \in \mathcal{C}_2$ has μ -measure zero, which yields the following

Corollary 4.3 *Let $v_1 \in \mathcal{C}_1$ and $v_2 \in \mathcal{C}_2$, then the set $\{v_1 = v_2\} \cap \{v_1' \neq v_2'\}$ is of μ -measure zero.*

As a consequence of Corollary 4.3, we have that the aggregate income of the firms is independent of the TBR, namely

$$\sum_i I_i(C_1, C_2, f) = \int_{\Theta_0} (\theta v(\theta) - v'(\theta)) \mu(d\theta) + \sum_i \int_{\Theta_i} (\theta v_i'(\theta) - v_i(\theta)) \mu(d\theta).$$

Finally we define

$$A(C_1, C_2, f) \triangleq \sum_i A_i(C_1, C_2, f).$$

4.2 Socially efficient allocations

A market situation (C_1^*, C_2^*) , together with a TBR f^* is said to be a *socially efficient allocation*¹⁰ if it minimizes the aggregate risk in the economy and if it is individually rational at the agencies' level; in other words if $A(C_1^*, C_2^*, f^*)$ solves the problem

$$\Psi \triangleq \inf_v \inf_f \inf_X \{A(C_1, C_2, f) \mid -\text{Var}[X_i(\theta)] = v_i'(\theta), A_i(C_1, C_2, f) \leq \varrho_i(W_i)\}.$$

As it was the case in Lemma 2.2, the variance constraints $-\text{Var}[X_i(\theta)] = v_i'(\theta)$ capture the incentive compatibility of the catalogues C_i , while $A_i(C_1, C_2, f) \leq \varrho_i(W_i)$ is firm i 's individual rationality constraint. As we mentioned in Section 2.5, in absence of a competitive market one cannot rely on equilibrium pricing to take care of the issue of individual rationality. It is therefore necessary to work under the constraints $A_i(C_1, C_2, f) \leq \varrho_i(W_i)$, for otherwise one could end up with allocations that are optimal on the aggregate level, yet unenforceable. An alternative would be to establish a cash-transfer system, which should be supervised by the regulator. We comment further on the latter in Section 4.2.4. In the remainder of this section we study the existence of SEAs, and we show that the presence a regulator is required in order for such allocations to be attainable and/or implementable. Our existence result depends heavily on the implementation of efficient TBRs.

Definition 4.4 *Let $C = (C_1, C_2)$ be a catalogue profile. A tie-breaking rule $\bar{f} \in F$ is **efficient** for C if*

$$A(C_1, C_2, \bar{f}) = \inf_{f \in F} \{A(C_1, C_2, f)\}.$$

From the definition above one observes that efficient TBRs are endogenously determined. Here we encounter the first need for the social planner in our risk-minimization setting, as it could be the case that unless regulated, efficient TBRs would not be implemented.

¹⁰From this point on we use the shorthand SEA to refer to a socially efficient allocation.

4.2.1 Minimizing the risk for fixed incomes and tie-breaking rule

In a first step, we fix $(v_1, v_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ (hence the firms' incomes), as well as $f \in F$. We shall abuse notation slightly and write $R_i(v_1, v_2, X_i, f)$ for $R_i(\mathcal{C}_1, \mathcal{C}_2, f)$. We then analyze the problem

$$\Psi_1 \triangleq \inf_{(X_1, X_2)} \left\{ \sum_i R_i(v_1, v_2, X_i, f) \mid -\text{Var}[X_i(\theta)] = v'_i(\theta) \text{ and } A_i(v_1, v_2, X_i, f) \leq \varrho_i(W_i) \right\}.$$

This problem can be decoupled into the sum of the infima, since the choice of X_i bears no weight on the evaluation of R_{-i} . Hence we must study the solution(s) to the following single-firm problems:

$$\inf_{X_i} \left\{ R_i(v_i, v_{-i}, X_i, f) \mid -\text{Var}[X_i(\theta)] = v'_i(\theta) \text{ and } A_i(v_1, v_2, X_i, f) \leq \varrho_i(W_i) \right\}.$$

To deal with the individual rationality constraints, we define

$$\mathcal{X}_i^{v,f} \triangleq \{X_i \in \mathcal{X}_i \mid A_i(v_1, v_2, X_i, f) \leq \varrho_i(W_i)\},$$

which is the set of individually rational products for the given price schedule (v_1, v_2) , and

$$\mathcal{X}_i^{k_i^v} \triangleq \{X \in \mathcal{X}_i \mid \|X_i\|_2^2 \leq k_i^v\}$$

where $k_i^v \triangleq v_i(a) - v_i(0)$. The set $\mathcal{X}_i^{k_i^v}$ contains all the products that can be structured as to construct incentive compatible catalogues given (v_1, v_2) , with k_i^v providing a bound to the \mathbb{L}^2 -norm of the products via the constraint $-\text{Var}[X_i(\theta)] = v'_i(\theta)$. We elaborate further into the k_i^v 's in Appendix A, where we also provide an outline of the proof of existence of solutions to Ψ_1 stated in Lemma 4.5. The proof is analogous to that of Theorem 2.3 in [20], and we include it for completeness.

Lemma 4.5 *For $v \in \mathcal{C}_1 \times \mathcal{C}_2$ and $f \in F$ given, if $\mathcal{X}_i^{k_i^v} \cap \mathcal{X}_i^{v,f} \neq \emptyset$ then there exist $X_i^{v,f} \in \mathcal{X}_i$ such that*

$$\Psi_1 = R_1(v_1, v_2, X_1^{v,f}, f) + R_2(v_1, v_2, X_2^{v,f}, f).$$

Remark 4.6 *It should be mentioned that individual rationality may be verified ex-post. If a solution $X_i^{v,f}$ to the minimization problem*

$$\inf_{X_i} \left\{ R_i(v_i, v_{-i}, X_i, f) \mid -\text{Var}[X_i(\theta)] = v'_i(\theta) \right\}$$

satisfies $A_i(v_1, v_2, X_i^{v,f}, f) \leq \varrho_i(W_i)$, then $\mathcal{X}_i^{k_i^v} \cap \mathcal{X}_i^{v,f} \neq \emptyset$. If on the contrary, $A_i(v_1, v_2, X_i^{v,f}, f) > \varrho_i(W_i)$, then the solution set of Ψ_1 is empty.

In order to establish the existence of an efficient TBR it is important to characterize the optimal contracts $X_i^{v,f}$. Specifically, we show below that $X_i^{v,f}$ can be multiplicatively decomposed into a θ -dependent function and an ω -dependent random variable. To this end, we construct the Lagrangian associated to minimizing $R_i(v_i, v_{-i}, X_i, f)$ subject to the moment conditions $E[X_i(\theta)] = 0$ and $\text{Var}[X_i(\theta)] + v'_i(\theta) = 0$, and we compute the Fréchet differential of R_i at X_i in the direction of $h \in \mathcal{X}_i$:

$$R'_i(X_i)h = \varrho'_i \left(W_i - \int_{\Theta_i} X_i(\theta) \mu(d\theta) - \int_{\Theta_0} X_i(\theta) f_i(\theta) \mu(d\theta) \right) \left(- \int_{\Theta_i} h(\theta) d\theta - \int_{\Theta_0} h(\theta) f_i(\theta) \mu(d\theta) \right).$$

Since, for all $H \in L^2(\Omega, \mathbb{P})$, the map $K \mapsto \varrho'_i(H)K$ is linear, it follows from the Riesz representation theorem that there is a random variable $Z_{X_i^{v,f}} \in L^2(\Omega, \mathbb{P})$ such that

$$R'_i(X_i)h = \int_{\Omega} Z_{X_i^{v,f}} \left(- \int_{\Theta_i} h(\theta) d\theta - \int_{\Theta_0} h(\theta) f_i(\theta) \mu(d\theta) \right) d\mathbb{P}.$$

Let $g : F \rightarrow \mathbb{L}^0(\Theta, \mu)$ be given by

$$g(f) = \mathbb{1}_{\Theta_i} + \mathbb{1}_{\Theta_0} f.$$

The operator R_i has an extremum at X_i under our moment constraints if there exist Lagrange multipliers $\lambda_i, \eta_i \in \mathbb{L}^2(\Theta, \mu)$ such that

$$\int_{\Omega} \int_{\Theta} h(\theta) \left(-Z_{X_i^{v,f}} g(f_i)(\theta) + \eta_i(\theta) + 2\lambda_i(\theta) X_i(\theta) \right) \mu(d\theta) d\mathbb{P} = 0.$$

for all $h \in \mathcal{X}_i$. Since $(\theta, \omega) \rightarrow -Z_{X_i^{v,f}} g(f_i)(\theta) + \eta_i(\theta) + 2\lambda_i(\theta) X_i(\theta)$ is an integrable function, the DuBois–Reymond lemma implies

$$-Z_{X_i^{v,f}} g(f_i)(\theta) + \eta_i(\theta) + 2\lambda_i(\theta) X_i(\theta) = 0. \quad (5)$$

Using the moment conditions $\mathbb{E}_{\mathbb{P}}[X_i^{v,f}(\theta)] = 0$ and $\text{Var}[X_i^{v,f}(\theta)] = -v'_i(\theta)$ we obtain

$$\eta_i(\theta) = \mathbb{E}_{\mathbb{P}}[Z_{X_i^{v,f}}] g(f_i)(\theta) \quad \text{and} \quad \lambda_i(\theta) = \frac{g(f_i)(\theta) \sqrt{\text{Var}[Z_{X_i^{v,f}}]}}{2\sqrt{-v'_i(\theta)}}.$$

Inserting the expressions for the Lagrange multipliers into Equation (5) yields

$$X_i^{v,f}(\theta) = \sqrt{-v'_i(\theta)} \bar{Z}_i^{v,f}, \quad (6)$$

where

$$\bar{Z}_i^{v,f} \triangleq \left(Z_{X_i^{v,f}} - \mathbb{E}[Z_{X_i^{v,f}}] \right) / \sqrt{\text{Var}[Z_{X_i^{v,f}}]}.$$

Equation (6) shows that the minimizers of problem Ψ_1 form collinear families in \mathcal{X} . Moreover, the randomness stemming from (Ω, \mathbb{P}) and the one induced by (Θ, μ) are decoupled. This property will prove to be key in Section 4.2.2, where we show the existence of efficient TBRs. The previous discussion yields the following characterization result:

Proposition 4.7 *Let $(X_1^{v,f}, X_2^{v,f})$ be a solution to Ψ_1 . Then $X_i^{v,f}$ takes the form*

$$X_i^{v,f}(\theta, \omega) = \sqrt{-v'_i(\theta)} \bar{Z}_i^{v,f}(\omega)$$

for some normalized random variable $\bar{Z}_i^{v,f}$ on $(\Omega, \mathcal{F}, \mathbb{P})$.

If either $X_1^{v,f}$ or $X_2^{v,f}$ do not satisfy $A_i(v_1, v_2, X_i^{v,f}, f) \leq \varrho_i(W_i)$, then using the standard convention $\inf\{\emptyset\} = \infty$ we would have $\Psi_1 = \infty$. In terms of the program Ψ , this guarantees that only pricing schedules and TBRs that offer the possibility of constructing individually rational catalogues stand a chance to be chosen.

4.2.2 Existence of efficient tie-breaking rules

In this section we show that for a given price schedule (v_1, v_2) there is a TBR f^v such that the corresponding optimal product lines $(X_1^{v,f^v}, X_2^{v,f^v})$ minimize the aggregate risk evaluation of the firms. For $v = (v_1, v_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ we define

$$\tilde{R}_i(v_1, v_2, f) \triangleq \varrho_i \left(W_i - Z_i^{v,f}(\omega) \left(\int_{\Theta_i} \sqrt{-v'_i(\theta)} d\theta + \int_{\Theta_0} f_i(\theta) \sqrt{-v'_i(\theta)} \mu(d\theta) \right) \right)$$

and

$$F_v \triangleq \left\{ f \in F \mid \sum_i \tilde{R}_i(v_1, v_2, f) < \infty \right\}.$$

We then have to solve the problem

$$\Psi_2 \triangleq \inf_{F_v} \sum_i \tilde{R}_i(v_1, v_2, f).$$

If $F_v = \emptyset$, then as above $\Psi_2 = \infty$. Otherwise we must verify the lower semicontinuity of the mapping

$$f \mapsto \tilde{R}_1(v_1, v_2, f) + \tilde{R}_2(v_1, v_2, f), \quad f \in F_v.$$

To this end consider a minimizing sequence $\{f^n\} \subset F_v$. The Banach–Alaoglu theorem guarantees that $\{f^n\}$ is weakly convergent¹¹ to some $f^v \in F_v$. Since $V_i(\theta) \triangleq \sqrt{-v'_i(\theta)}$ belongs to $\mathbb{L}^2(\Theta, \mu)$ for all $v_i \in \mathcal{C}_i$, then

$$\int_{\Theta_0} f_i^n(\theta) V_i(\theta) d\theta \rightarrow \int_{\Theta_0} f_i^v(\theta) V_i(\theta) \mu(d\theta). \quad (7)$$

Let

$$a_i^{v,n} \triangleq \int_{\Theta_i} V_i(\theta) d\theta + \int_{\Theta_0} f_i^n(\theta) V_i(\theta) \mu(d\theta).$$

For f^v we define a_i^v analogously. From (7) we have that $a_i^{v,n} \rightarrow a_i^v$. Since $\|Z_i^{v,f^n}\|_2 \leq 1$ for all n , there exists $Z_i^v \in \mathbb{L}^2(\Omega, \mathbb{P})$ such that $\|Z_i^v\|_2 \leq 1$ and

$$Z_i^{v,f^n} \xrightarrow{w.} Z_i^v.$$

Therefore $Y_i^{v,n} \triangleq a_i^{v,n} \cdot Z_i^{v,f^n}$ converges weakly to $Y_n^v \triangleq a_i^v \cdot Z_i^v$. It follows from the Fatou property of ϱ_i that

$$\varrho_i(W_i - a_i^v \cdot Z_i^v) \leq \liminf_{n \rightarrow \infty} \varrho_i(W_i - a_i^{v,n} \cdot Z_i^{v,f^n}).$$

From the lower semicontinuity of the norm in terms of weak convergence we have

$$\text{Var}[Z_i^v V_i(\theta)] \leq -v'_i(\theta).$$

Proceeding as in Section 4.2.1, we have that

$$\tilde{R}_i(v_1, v_2, f^v) \leq \varrho_i(W_i - a_i^v \cdot Z_i^v).$$

Therefore

$$\sum_{i=1}^2 (\tilde{R}_i(v_1, v_2, f^v) - I_i(v_1, v_2, f^v)) = \inf_{f \in F} \left\{ \sum_{i=1}^2 (\tilde{R}_i(v_1, v_2, f) - I_i(v_1, v_2, f)) \right\}.$$

We denote by (X_1^v, X_2^v) any optimal list of claims associated to the pricing schedules (v_1, v_2) and the TBR f^v . For notational convenience we define

$$A(v_1, v_2) \triangleq \sum_{i=1}^2 (\tilde{R}_i(v_1, v_2, f^v) - I_i(v_1, v_2, f^v)).$$

4.2.3 Minimizing with respect to the firms' incomes

To finalize the proof of existence of a SEA, we let $\{(v_1^n, v_2^n)\}$ be a minimizing sequence of

$$\sum_{i=1}^2 (\tilde{R}_i(v_1, v_2, f^v) - I_i(v_1, v_2, f^v)).$$

¹¹We omit writing “up to a subsequence” when exploiting the compactness of sets, but it should of course be kept in mind.

We get from Proposition 2.3 that there exist $(\bar{v}_1, \bar{v}_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ such that $v_i^n \rightarrow \bar{v}_i$ uniformly. This implies that $V_i^n \rightarrow \bar{V}_i$ almost surely (see for example Proposition A.4 in [11]). Moreover, for any $\theta \in \Theta$ where convergence holds, it is uniform. We have from Fatou's lemma and the fact that $0 \leq -I_i(v_1, v_2, f)$ for any $(v_1, v_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ and any $f \in F$ that

$$\sum_{i=1}^2 -I_i(\bar{v}_1, \bar{v}_2, f^{\bar{v}}) \leq \liminf_{n \rightarrow \infty} \sum_{i=1}^2 -I_i(v_1^n, v_2^n, f^{v^n}).$$

In order to deal with the risky part of the firms' problems we require the following

Lemma 4.8 *Let $\{\phi_n\}, \{\psi_n\} \subset \mathbb{L}^2(\Omega, \mathbb{P})$ and $\phi, \psi \in \mathbb{L}^2(\Omega, \mathbb{P})$ be such that*

$$\phi_n \xrightarrow{\|\cdot\|_2} \phi \quad \text{and} \quad \psi_n \xrightarrow{w.} \psi, \quad \text{then} \quad \langle \phi_n, \psi_n \rangle \rightarrow \langle \phi, \psi \rangle,$$

where $\langle \cdot, \cdot \rangle$ is the canonical inner product in $\mathbb{L}^2(\Omega, \mathbb{P})$.

PROOF. Adding and subtracting $\langle \phi, \psi_n \rangle$ from $|\langle \phi_n, \psi_n \rangle - \langle \phi, \psi \rangle|$ we obtain

$$|\langle \phi_n, \psi_n \rangle - \langle \phi, \psi \rangle| \leq |\langle \phi_n - \phi, \psi_n \rangle| + |\langle \phi, \psi_n - \psi \rangle|.$$

Since $\{\psi_n\}$ is a weakly convergent sequence, it is bounded. Let \bar{K} be such bound, then using the Cauchy-Schwarz inequality we have

$$|\langle \phi_n, \psi_n \rangle - \langle \phi, \psi \rangle| \leq \bar{K} \|\phi_n - \phi\|_2 + |\langle \phi, \psi_n - \psi \rangle|.$$

As $n \rightarrow \infty$, the first summand on the righthand side of the inequality converges to zero due to the strong convergence of the ϕ_n 's to ϕ ; the second summand converges to zero due to the weak convergence of the ψ_n 's to ψ , which concludes the proof. \square

We show in Proposition 4.9 that if it were the case that the market segments exhibited no jumps in their limiting behavior then there would exist SEAs. By absence of jumps we mean that

$$\mathbb{I}_{\Theta_0^n} \xrightarrow{a.s.} \mathbb{I}_{\bar{\Theta}_0} \quad \text{and} \quad \mathbb{I}_{\Theta_i^n} \xrightarrow{a.s.} \mathbb{I}_{\bar{\Theta}_i},$$

where $\bar{\Theta}_0 = \{\bar{v}_1 = \bar{v}_2\}$ and $\bar{\Theta}_i = \{\bar{v}_i > \bar{v}_{-i}\}$. This ‘‘nice’’ convergence of the Θ_i^n 's is by no means the general scenario. As an example let $v_1^n \equiv v_1$ and $v_2^n = v_1 + 1/n$. Here $\Theta_1^n \equiv \Theta$, $\Theta_2^n \equiv \emptyset$, but in the limit only $\bar{\Theta}_0 \neq \emptyset$. Nonetheless, the existence result of SEAs under the assumption of convergence of the indicator functions is relevant for the general case, and we present it below.

Proposition 4.9 *Let $\{(v_1^n, v_2^n)\}, (\bar{v}_1, \bar{v}_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ such that $v_i^n \rightarrow \bar{v}_i$ uniformly. Assume that $\mathbb{I}_{\Theta_j^n} \rightarrow \mathbb{I}_{\bar{\Theta}_j}$, and let \bar{X}_1, \bar{X}_2 and \bar{f} be the (weak) limits of $X_1^{v^n}, X_2^{v^n}$ and f^{v^n} . Then*

$$\begin{aligned} \liminf_{n \rightarrow \infty} \sum_{i=1}^2 \tilde{R}_i(v_1^n, v_2^n, f^{v^n}) &\geq \sum_{i=1}^2 \varrho_i \left(W_i - \int_{\bar{\Theta}_i} \bar{X}_i \mu(d\theta) - \int_{\bar{\Theta}_0} \bar{f}_i \bar{X}_i \mu(d\theta) \right) \\ &\geq \sum_{i=1}^2 \varrho_i \left(W_i - \int_{\bar{\Theta}_i} X_i^{\bar{v}} \mu(d\theta) - \int_{\bar{\Theta}_0} \bar{f}_i^{\bar{v}} X_i^{\bar{v}} \mu(d\theta) \right), \end{aligned}$$

where $(X_1^{\bar{v}}, X_2^{\bar{v}}, f^{\bar{v}})$ solves the social planner's problem for (\bar{v}_1, \bar{v}_2) .

PROOF. The second inequality follows from the definition of $(\bar{X}_1, \bar{X}_2, \bar{f})$. To show that the first one holds, we first use Proposition 4.7 and write

$$\int_{\Theta_i^n} X_i^{v^n, f^{v^n}} \mu(d\theta) - \int_{\Theta_0^n} f_i^n X_i^{v^n, f^{v^n}} \mu(d\theta) = Z_i^{v^n, f^{v^n}} \int_{\Theta} \left(\mathbb{1}_{\Theta_i^n}(\theta) - f_i^{v^n} \mathbb{1}_{\Theta_0^n} V_i^n(\theta) \right) \mu(d\theta) \triangleq a_i^{v^n} \cdot Z_i^{v^n, f^{v^n}}.$$

The assumption on the convergence of the indicator functions $\mathbb{1}_{\Theta_j^n}$ implies that

$$V_i^n \mathbb{1}_{\Theta_0^n} \xrightarrow{a.s.} \bar{V}_i \mathbb{1}_{\bar{\Theta}_0}.$$

Since $|V_i^n \mathbb{1}_{\Theta_0^n}|, |\bar{V}_i \mathbb{1}_{\bar{\Theta}_0}| \leq M$ and $\mu(\Omega) < \infty$, by Lebesgue Dominated Convergence we get

$$V_i^n \mathbb{1}_{\Theta_0^n} \xrightarrow{\|\cdot\|_2} \bar{V}_i \mathbb{1}_{\bar{\Theta}_0}.$$

Hence, by Lemma 4.8 we have that $a_i^{v^n} \rightarrow \bar{a}_i$, where

$$\bar{a}_i \triangleq \int_{\Theta} \left(\mathbb{1}_{\bar{\Theta}_1}(\theta) - \bar{f} \mathbb{1}_{\bar{\Theta}_0} \bar{V}_i(\theta) \right) \mu(d\theta).$$

Therefore

$$a_i^{v^n} \cdot Z_i^{v^n} \xrightarrow{w.} \bar{a}_i \cdot \bar{Z}_i,$$

and the Fatou property of the risk measure yields

$$\sum_{i=1}^2 R_i(\bar{v}_1, \bar{v}_2, X_i, f) \leq \liminf_{n \rightarrow \infty} \sum_{i=1}^2 \tilde{R}_i(v_1^n, v_2^n, f^{v^n}).$$

□

We now deal with the possibility of non-convergence of the indicator functions. We first observe that the Θ_0^n 's are closed subsets of Θ , which is compact. The set

$$\bar{2}^\Theta := \{A \subset \Theta \mid A \neq \emptyset \text{ and } A \text{ is closed}\}$$

endowed with the Hausdorff metric $h : \bar{2}^\Theta \rightarrow \mathbb{R}_+$ is a compact metric space (see for example [1], Chapter 3). If $\Theta_0^n \neq \emptyset$ infinitely often, then there exists $\hat{\Theta}_0 \in \bar{2}^\Theta$ such that, up to a subsequence if necessary

$$\Theta_0^n \xrightarrow{h} \hat{\Theta}_0.$$

Moreover, $\hat{\Theta}_0$ is contained in $\bar{\Theta}_0$. If, on the contrary, $\Theta_0^n = \emptyset$ for all but a finite number of n 's, then we define the Hausdorff limit of $\{\Theta_0^n\}$ as the empty set, which is again contained in $\bar{\Theta}_0$. In both instances we observe that there can only be more ties in the limit, and hence more ways of breaking them. This suggests that the aggregate risk in the limit is indeed no greater than the aggregate risk in the pre-limit.

As for the sets Θ_i^n , if θ is not eventually in either of them, then either $\theta \in \hat{\Theta}_0$ or it is a type whose preferences alternate. However, the limiting behavior of the latter is that of indifference, due to the convergence of the functions v_i^n . In other words, an agent type θ eventually always contracts with the same firm i , alternates between firms or is always indifferent. Thus,

$$\theta \in \Theta_0^n \text{ for all sufficiently large } n \text{ or } \theta \in \bar{\Theta}_0.$$

We now decompose the type space into subsets for which the associated indicator functions converge. To this end, we define

$$\tilde{\Theta}_0^n \triangleq \bar{\Theta}_0 \setminus \Theta_0^n, \quad \text{and} \quad \tilde{\Theta}_i^n \triangleq \Theta_i^n \setminus \tilde{\Theta}_0^n.$$

The sets $\tilde{\Theta}_i^n$ contain the “surviving customers”, in the sense that there will be no jumps towards indifference from types in $\tilde{\Theta}_i^n$ at the limit; $\tilde{\Theta}_0^n$ is the set of “alternating customers”. By construction

$$\Theta = \tilde{\Theta}_0^n \cup \Theta_0^n \cup \tilde{\Theta}_1^n \cup \tilde{\Theta}_2^n = \bar{\Theta}_0 \cup \bar{\Theta}_1 \cup \bar{\Theta}_2^n.$$

The following lemma shows that the sets $\tilde{\Theta}_i^n$ of “surviving customers” converge to the sets of agent types that contract with firm i when (\bar{v}_1, \bar{v}_2) is offered.

Lemma 4.10 *For $i = 1, 2$ we have $\mathbb{1}_{\tilde{\Theta}_i^n} \rightarrow \mathbb{1}_{\bar{\Theta}_i}$.*

PROOF. Let $\theta \in \Theta$ be given. There are two possible cases. Either there exists $N \in \mathbb{N}$ such that if $n > N$, then $\theta \in \tilde{\Theta}_i^n$ or there exists a subsequence $\{v_1^{n_k}, v_2^{n_k}\}$ such that $\theta \notin \tilde{\Theta}_i^{n_k}$ for all n_k . In the former case the (uniform) convergence of the v_i^n 's implies $\bar{v}_i(\theta) > \bar{v}_{-i}(\theta)$. In the latter case, either θ eventually belongs to $\tilde{\Theta}_{-i}^n$, in which case convergence of the indicator functions follows or $\theta \in \tilde{\Theta}_1^n \cap \tilde{\Theta}_2^n$ infinitely often. In this case $\theta \in \tilde{\Theta}_0^n$ and we again have convergence of the indicators. \square

Let us now assume that the social planner's problem were such that at every stage n all agent types that belong to the set of “alternating customers” are deemed indifferent. This results in more tie-breaking possibilities. Specifically we consider the social planner's problem

$$A^*(v_1^n, v_2^n) = \inf_{f^n \in F} \inf_X \left\{ \sum_{i=1}^2 \varrho_i \left(W_i - \int_{\tilde{\Theta}_i^n} X_i^n \mu(d\theta) - \int_{\tilde{\Theta}_0^n \cup \Theta_0^n} f_i^n X_i^n \mu(d\theta) \right) \middle| \text{Var}[X_i^n(\theta)] = -(v_i^n)'(\theta) \right\}.$$

A subset of the possible choices of the f_i^n 's is F_0^n , which is defined as the set of $\tilde{f} \in F$ such that

$$\tilde{f}(\theta) = \begin{cases} 0, & v_i^n(\theta) < v_{-i}^n(\theta) \\ 1, & v_i^n(\theta) > v_{-i}^n(\theta) \\ f(\theta), & \text{otherwise,} \end{cases}$$

for some TBR $f \in F$. We then have that

$$A^*(v_1^n, v_2^n)|_{F_0} = A(v_1^n, v_2^n).$$

Thus,

$$A(v_1^n, v_2^n) \geq A^*(v_1^n, v_2^n).$$

Consider a minimizing sequence $\{(v_1^n, v_2^n)\}$ with (uniform) limit (\bar{v}_1, \bar{v}_2) . Lemma 4.10 guarantees that $\mathbb{1}_{\tilde{\Theta}_i^n} \rightarrow \mathbb{1}_{\bar{\Theta}_i}$; moreover $\tilde{\Theta}_0^n \cup \Theta_0^n \equiv \bar{\Theta}_0$ so patently $\mathbb{1}_{\tilde{\Theta}_0^n \cup \Theta_0^n} \rightarrow \mathbb{1}_{\bar{\Theta}_0}$. These two facts allow us to apply Proposition 4.9 to $A^*(v_1^n, v_2^n)$, which yields

$$\begin{aligned} \liminf_{n \rightarrow \infty} A(v_1^n, v_2^n) &\geq \liminf_{n \rightarrow \infty} A^*(v_1^n, v_2^n) \\ &\geq \sum_{i=1}^2 \varrho_i \left(W_i - \int_{\bar{\Theta}_i} X_i^* \mu(d\theta) - \int_{\bar{\Theta}_0} f_i^* X_i^* \mu(d\theta) \right), \end{aligned}$$

where (X_1^*, X_2^*, f^*) solves the social planner's problem for (\bar{v}_1, \bar{v}_2) . This shows that (X_1^*, X_2^*, f^*) is indeed optimal because $\{(v_1^n, v_2^n)\}$ was required to be a minimizing sequence. We have proved the following

Theorem 4.11 *If firm i assesses risk using a law invariant risk measure*

$$\varrho_i : \mathcal{X} \rightarrow \mathbb{R} \cup \{+\infty\},$$

which has the Fatou property, and if it offers catalogues of the form

$$C_i = \{(X_i(\theta), p_i(\theta))\}_{\theta \in \Theta},$$

then there exists a socially efficient market situation.

4.2.4 Individual rationality revisited

We can also deal with a slightly different setting, in which the regulator’s presence would be necessary for the enforcement of a cash transfer. Namely, assume that (C_1^*, C_2^*, f^*) is a solution to the problem Ψ without the individual rationality constraints $\hat{R}_i(C_1, C_2, f) \leq \varrho_i(W_i)$. Since both firms could simply offer $(0, 0)$, a minimization of $A(\cdot, \cdot, \cdot)$ would be IR on the aggregate level. Then there would exist $r \geq 0$ (the *rent of risk exchange*) such that

$$A(C_1^*, C_2^*, f^*) = \varrho_1(W_1) + \varrho_2(W_2) - r$$

We define a *transfer SEA* to be a quadruple (C_1^*, C_2^*, f^*, T^*) such that

1. $\Psi = A(C_1^*, C_2^*, f^*)$,
2. $\hat{R}_1(C_1^*, C_2^*, f^*) - T^* \leq \varrho_1(W_1)$,
3. $\hat{R}_2(C_1^*, C_2^*, f^*) + T^* \leq \varrho_2(W_2)$,

and $T^* \in [r_1, r_2]$ for some $r_1, r_2 \in \mathbb{R}$ such that $r_2 - r_1 = r$. The arguments contained in Sections 4.2.1, 4.2.2 and 4.2.3 can be immediately applied to prove the following

Corollary 4.12 *If firm i assesses risk using a law invariant risk measure*

$$\varrho_i : \mathcal{X} \rightarrow \mathbb{R} \cup \{+\infty\},$$

which has the Fatou property, and if it offers catalogues of the form

$$C_i = \{(X_i(\theta), p_i(\theta))\}_{\theta \in \Theta},$$

then there exists a transfer socially efficient market situation.

5 Examples: Risk Minimization

In this section we focus our attention on two well-known risk measures: entropic and average value at risk. When it comes to the latter, we mostly present in Section 5.3 the results of applying the numerical algorithm that is described in Appendix B to some specific examples¹². As for the entropic risk measure, before proceeding with the numerical simulations, we present in Section 5.1 a structural result that exploits the risk measure’s particular structure. Moreover, we show that in this case there is a SEA where both firms service all of the agents, and for which the optimal TBR is a constant proportion over the whole market. We refer to such efficient TBRs as “fix-mix” rules.

5.1 Entropic-risk-minimizing firms

In what follows we concentrate on a situation where the firms use the entropic risk measure as a means to assess their risk exposure, i.e.

$$\varrho_i(X) = \frac{1}{\gamma_i} \ln \mathbb{E}_{\mathbb{P}} [\exp(-\gamma_i X)].$$

The coefficient γ_i represents firm i ’s risk aversion. This particular choice of risk measures, which is closely related to exponential utility, allows us to further the analysis into the structure of SEAs. Moreover, given

¹²All of our codes are available upon request

that it is strictly convex and that it has a closed-form representation it lends itself very well to numerical exercises. For v_i and f given, Proposition 4.7 allows us to write program Ψ_1 (the minimization with respect to the firms' claims for fixed incomes and TBR in Section 4.2.1) as

$$\inf_{(Z_1, Z_2) \in \Gamma_i \times \Gamma_2} \sum_{i=1}^2 \varrho \left(W - Z_i \left(\int_{\Theta_i} \sqrt{-v'_i(\theta)} \mu(d\theta) + \int_{\Theta_0} \sqrt{-v'_i(\theta)} f_i(\theta) \mu(d\theta) \right) \right),$$

where

$$\Gamma_i := \{Z \in \mathcal{X}_i \mid \mathbb{E}[Z] = 0, \|Z\|_2^2 = 1\}.$$

The structure above allows us to write (See Section 3.4.2 in [20]) the minimization problem of firm i as that of finding a stationary point to the Lagrangian

$$L(Z, \tau, \kappa) = \ln \left(\int_{\Omega} e^{-\gamma_i(W + Z a(v_i, f_i))} d\mathbb{P} \right) + \tau \int_{\Omega} Z d\mathbb{P} + \kappa \int_{\Omega} Z^2 d\mathbb{P},$$

where τ and κ are the Lagrange multipliers associated to the moment constraints, and

$$a(v_i, f_i) = \int_{\Theta_i} \sqrt{-v'_i(\theta)} \mu(d\theta) + \int_{\Theta_0} \sqrt{-v'_i(\theta)} f_i(\theta) \mu(d\theta).$$

This in turn results in the following implicit representation for the optimal claims given v_i and f , which has a unique solution for each realization $z_i^{v, f}$ of $Z_i^{v, f}$ and w_i of W_i :

$$Z_i^{v, f} = - \frac{e^{-\gamma_i(W_i + Z_i^{v, f} a(v_i, f_i))} - \mathbb{E} \left[e^{-\gamma_i(W_i + Z_i^{v, f} a(v_i, f_i))} \right]}{\sqrt{\text{Var} \left(e^{-\gamma_i(W_i + Z_i^{v, f} a(v_i, f_i))} \right)}}. \quad (8)$$

Remark 5.1 *The indirect utility functions v_i and the TBR f_i only affect the optimal claims $Z_i^{v, f}$ via the “aggregator” $a(v_i, f_i)$.*

The previous remark has interesting repercussions for the socially efficient TBRs. Indeed, assume the firms have initial endowments W_i and they are characterized by risk aversion coefficients γ_i . Theorem 4.11 guarantees the existence of a SEA $(v_1^*, v_2^*, f^*, Z_1^*, Z_2^*)$, which yields the following aggregate risk in the economy:

$$\varrho_1(W_1 - a(v_1^*, f_1^*) Z_1^*) + \varrho_2(W_2 - a(v_2^*, f_2^*) Z_2^*) + I[v_1^*, v_2^*], \quad (9)$$

where $I[v_1^*, v_2^*]$ represents the firms' aggregate income (which is independent of f^*). We know from Remark 5.1 that any modification on f^* that leaves $a(v_i^*, f_i^*)$ unchanged bears no weight on the value of expression (9). We can redefine

$$f^*(\theta) = \begin{cases} 0, & v_1(\theta) < v_2(\theta) \\ 1, & v_1(\theta) > v_2(\theta) \\ f^*(\theta), & \text{otherwise,} \end{cases}$$

and define

$$v^*(\theta) \triangleq \max \{v_1^*(\theta), v_2^*(\theta)\},$$

which then allows us to write

$$a(v_i^*, f_i^*) = \int_{\Theta} \sqrt{-(v^*)'(\theta)} f_i^*(\theta) \mu(d\theta).$$

Once we have defined the TBR over the whole Θ , we may go one step further and write

$$K \triangleq \frac{\int_{\Theta} \sqrt{-(v^*)'(\theta)} f^*(\theta) \mu(d\theta)}{\int_{\Theta} \sqrt{-(v^*)'(\theta)} \mu(d\theta)},$$

then

$$a(v_1^*, f_1^*) = K \int_{\Theta} \sqrt{-(v^*)'(\theta)} \mu(d\theta) \quad \text{and} \quad a(v_2^*, f_2^*) = (1 - K) \int_{\Theta} \sqrt{-(v^*)'(\theta)} \mu(d\theta).$$

An interesting economical conclusion from the computations above is that there exists a SEA that consists of both firms servicing the whole market, and splitting the customers in a K to $1 - K$ proportion. This follows from the fact that only v^* (the upper envelope of the original v_i^*) appears in each firm's program. In markets such as regulated health-insurance, where firms are legally prevented from abstaining from contracting with any agent, the regulator may oversee that a socially optimal proportion of the market is serviced by each firm.

5.2 Simulations (Entropic-risk minimizing firms)

In this section we provide the numerical analysis of a particular example of the entropic-risk minimizing firms analyzed above. To this end we set


- Dimension of the space defining v : 6, dimension of Ω : 14,
- $W_1 = 0.5 * (-1, -3, -9, -3, -1, -0.2, -0.1, -0.1, -0.2, 1, -3, -9, -3, -1)^T$,
- $W_2 = 0.5 * (-0.03, -0.1, -0.18, -0.2, -1, -3, -9, -10, -3, -1, -0.2, -0.18, -0.1, -0.03)^T$,
- risk aversion coefficient $\gamma = 2$.

5.2.1 Risk assessments


Let us benchmark the aggregate risk in the competitive economy against the risk in a monopolistic setting. To this end, we first fix $f \equiv 1$. This corresponds to a model in which firm 1 acts as a monopolist and firm 2 has no access to the market. The a-priori aggregate risk in the economy is 7.36. The risk assessment of firm 1 is reduced from 3.53 to 2.16 after the it has traded with the agents while the risk of firm 2 stays the same. Below we plot the numerical result for Z , as well as the theoretical one from Equation (8) in Figure 1(a) and the corresponding minimizing v (the agents' indirect utility) in Figure 1(b). Likewise, if we fix $f \equiv 0$, which again yields a principal-agent setting (competition-wise) for firm 2, we obtain that this firm's initial risk is 3.84, which is reduced to 2.16 after trading. Finally, once we let f vary, the aggregate risk decreases from 7.36 to 5.39, and the corresponding final risk assessments for the firms are 2.17 and 2.67 respectively. While each individual firm is worse off in the presence of competition, it should be noted that the decrement in the aggregate risk in the economy is lower in either case where only one firm has access to the market. Risk decreases from 7.36 to 5.99 when only firm 1 is active, and from 7.36 to 5.72 if it is firm 2 who trades with the agents.

5.2.2 Risk profiles

With respect to the ex-ante and ex-post risk profiles, we observe that trading has a smoothing effect, flattening spikes that correspond to the bad states of the World. However the basic shape of the risk profile remains, which is in contrast to what we find in Figure 4 of our AV@R example. This is presented in Figure 2, where the 14 elementary events have been connected by lines for illustration purposes. Figure 3




(a) Numerical v.s. theoretical values for Z




(b) The minimizing v

Figure 1: Comparison results for the $f \equiv 1$ case.



(a) Comparison of W_1 and $W_1 - a(v_1^*, f_1^*)Z_1^*$



(b) Comparison of W_2 and $W_2 - a(v_2^*, f_2^*)Z_2^*$

Figure 2: Positions before and after trading.

shows the agents' indirect utilities associated to each firm's offer, and compares it to the indirect utility for the agents who face a monopolist (in this case firm 1). The market is shared at a 0.42 to 0.58 ratio between the firms. We observe that the upper envelope $u(\theta) = \max\{v_1(\theta), v_2(\theta)\}$ dominates the monopolistic situation for all agents. In conclusion, this particular example is in line with the intuition that competition among sellers benefits the buyers. Moreover, the competitive setting also provides a lower aggregate risk exposure. A point could be made that the regulator should make sure that enough incentives exist for all firms to engage in risk-minimizing trading, as it is socially desirable.

Remark 5.2 *Arguably, the implementation of an efficient TBR in full generality could prove to be a daunting task. However, the structure of $Z_i^{v,f}$ presented in Equation (8) suggests that in general these variables depend on v and f only via the integral expression $a(v, f)$. If such were the case, one could achieve an optimum by choosing an efficient "fix-mix" TBR, where each firm caters to a constant proportion of*




Figure 3: Indirect utilities under monopoly and oligopoly.

the whole market. The latter clearly makes the implementation considerably simpler. For the special case of the entropic measure this can indeed be achieved, as shown above.

5.3 Simulations (AV@R–minimizing firms)

In this section we study an example where the firms are average value-at-risk minimizers. Recall that for $\lambda \in (0, 1]$, and $X \in \mathbb{L}^\infty(\Omega, \mathbb{F}, P)$, one defines

$$AV@R_\lambda(X) \triangleq \frac{1}{\lambda} \int_0^\lambda V@R_\gamma(X) d\gamma,$$


where

$$V@R_\gamma(X) \triangleq \inf \{m \mid P\{X + m < 0\} \leq \gamma\}.$$


We use the following initial parameters:

- Dimension of the space defining v : 6, dimension of (Ω) : 14,
- $W_1 = 0.02 * (-1, -2, -4, -10, -4, -2, -1, -0.8, -0.5, -0.3, 0, 0, 0, 0)^T$,
- $W_2 = 0.05 * (-0.03, -0.1, -0.18, -0.2, -1, -3, -9, -10, -3, -1, -0.2, -0.18, -0.1, -0.03)^T$,
- levels for the AV@R: λ_1 : 0.05 and λ_2 : 0.1

The initial aggregate initial risk assessment is 0.68, which decreases to 0.21 after trading. In Figure 4 we compare the firms' positions before and after trading. We observe that in contrast with the entropic-risk-measure case (see Figure 2), the ex-post shapes of the risk profiles have been significantly altered. This is due to the fact that the risk measure in hand places a heavier weight on the bad states of nature, even at the cost of the originally good ones Figure 5(a) shows the indirect utility functions corresponding to the catalogue that each firm offers, and in Figure 5(b) we have plotted an efficient TBR.




(a) Comparison of W_1 and $W_1 - a(v_1^*, f_1^*)Z_1^*$




(b) Comparison of W_2 and $W_2 - a(v_2^*, f_2^*)Z_2^*$

Figure 4: Positions before and after trading.



(a) The indirect utility functions



(b) A possible socially efficient TBR

Figure 5: Indirect utilities and a possible efficient TBR.

Remark 5.3 *It should be noted that in neither of the examples presented above were the individual rationality constraints for the firms binding. This is essential to implement the “fix-mix” TBR without the need of a cash transfer. Although f plays no role in the aggregate income $I[v_1, v_2]$, it does enter the IR constraints of each firm. If one wanted to implement a “fix-mix” TBR without a cash transfer scheme, they would in general run into a two-equations-one-unknown issue as soon as an IR constraint bound.*

6 Conclusions

In this paper we have extended the principal-agent models of profit maximization found in [9] and risk minimization found in [20] to a multi-agency setting. Both of these works deal with OTC trading of derivatives under adverse selection. It should be mentioned that in order to deal with the increased

complexity introduced by the enlargement of the set of sellers, additional restrictions had to be imposed on the set of financially feasible products. On the profit maximization side we made use of results of Reny and, more recently, Page Jr. & Monteiro to show the existence of (mixed–strategies) Nash equilibria. We were unable, however, to obtain a similar result in the case of risk minimization. The non–linear impact of individual contracts on the firms’ risk evaluations is incompatible with the machinery developed by the authors previously mentioned. Moreover, even if the game were uniformly payoff secure and (weakly) reciprocal upper semicontinuous, the lack of quasi–convexity would require considering the mixed extension of the game. This brings us to the following conceptual issue: what is expected risk? Convex risk measures can be (robustly) represented as a worst–case analysis over a family of probability measures that are absolutely continuous with respect to the reference one. However, not all of these measures (which can be interpreted as possible distributions of the future states of the World) are given the same weight. Indeed, they are penalized according to a function that maps the space of probability measures into the extended Reals. By considering a mixed extension of our risk–minimization game, we implicitly assume that “Nature” and the firms behave in qualitatively different ways. In our view this would be inconsistent. Instead, we introduce the notion of socially efficient allocations and prove the existence of such. We believe that an important contribution of this paper is to show that, within our stylized setting, non–regulated, OTC markets cannot be guaranteed to be efficient (in the sense of the welfare theorems). The extension of our general setting to one where agents have heterogenous initial endowments (multi–dimensional agent types) and (we believe more interestingly) to a dynamic framework are left for future research.

A Existence of minimizers to Ψ_1

In this appendix we give an overview of the proof of existence of minimizers to problem Ψ_1 in Section 4.2.1, which is analogous to the proof of Theorem 2.3 in [20]. The main ideas behind the proof are to relax the variance constraint to have a convex minimization problem, and then to show that based on a solution to the latter we can construct a solution to Ψ_1 . The steps to be taken are the following:

1. We fix $(v_1, v_2) \in \mathcal{C}_1 \times \mathcal{C}_2$ and relax the variance constraint to $-\text{Var}[X_i(\theta)] \leq v'_i(\theta)$. We observe that

$$\|X_i\|_2^2 = v_i(a) - v_i(0) \triangleq k_i^v.$$

Let K_i be the uniform bound on the elements of \mathcal{C}_i (see Prop. 2.3), then $\|X_i\|_2^2 \leq k_i^v \leq K_i$.

2. The convexity of ϱ_i implies that the set

$$\mathcal{X}_i^{v,f} \triangleq \{X_i \in \mathcal{X}_i \mid A_i(v_1, v_2, X_i, f) \leq \varrho_i(W_i)\}$$

is convex. The fact that ρ_i has the Fatou property implies that $\mathcal{X}_i^{v,f}$ is closed.

3. Since the mapping

$$X_i \mapsto \int_{\Theta_i} X_i(\theta) d\theta - \int_{\Theta_0} X_i(\theta) f_i(\theta) d\theta,$$

is a linear, then $X_i \mapsto A_i(v_1, v_2, X_i, f)$ is a convex mapping. The set

$$\mathcal{X}_i^{k_i^v} \triangleq \{X \in \mathcal{X}_i \mid \|X_i\|_2^2 \leq k_i^v\}$$

is a closed, convex and bounded set, hence the relaxed optimization problem has a solution $\tilde{X}_i^{v,f}$, as long as $\mathcal{X}_i^K \cap \mathcal{X}_i^{v,f} \neq \emptyset$.

4. The variance constraint can be made binding by structuring the products \tilde{X}_i^v into $X_i^{v,f} = \tilde{X}_i^{v,f} + \alpha Y$. Here Y is independent of θ and $\alpha : \Theta \rightarrow \mathbb{R}$ integrates to zero.

B Description of the algorithm used in Sections 5.2 and 5.3

In this appendix we provide a brief description of our numerical algorithm, whose aim is to estimate solutions to the problem

$$\inf_{(Z, f, v)} \sum_{i=1}^2 \left[\varrho_i(W_i - Z_i \int_a^1 \sqrt{-v'(\theta)} f_i(\theta) d\theta) + \int_a^1 (v(\theta) - \theta v'(\theta)) f_i(\theta) d\theta \right]$$

subject to:

- $f_i \geq 0$, $f_1 + f_2 = 1$;
- v convex, $v \geq 0$, $v' \leq 0$, $v(1) = 0$;
- $E[Z_i] = 0$, $\text{Var}[Z_i] = 1$.
- $\varrho_i(W_i - Z_i \int_a^1 \sqrt{-v'(\theta)} f_i(\theta) d\theta) + \int_a^1 (v(\theta) - \theta v'(\theta)) f_i(\theta) d\theta \leq \varrho_i(W_i)$. It should be noted that in the examples presented in Section 5, this constraint was not binding.

In order to do so, we set a discretization level n , and we work with the following structures:

- $f_1(\lambda) \triangleq \sum_{k=0}^{n-1} \lambda_k \mathbb{1}_{[a+k\frac{1-a}{n}, a+(k+1)\frac{1-a}{n}]}$, $f_2 = 1 - f_1$, $0 \leq \lambda_i \leq 1$.
- $v'(\cdot) \triangleq \left(\int_a^1 -\frac{d}{d\theta} \sqrt{-v'(\theta)} d\theta + \sqrt{-v'(1)} \right)^2$, we shall denote $\tilde{v} \triangleq -\frac{d}{d\theta} \sqrt{-v'(\theta)}$, and $\gamma \triangleq \sqrt{-v'(1)} \geq 0$. There is a one-to-one correspondence between $\tilde{v} \geq 0$, $\gamma \geq 0$ and functions v that satisfy the requirements of being convex, non-increasing and non-negative. This allows for an easier-to-handle description of such functions, instead of having to impose more burdensome, global convexity constraints, which would be necessary in higher dimensions.
- We go one step further and define $\tilde{v}(\alpha) \triangleq \sum_{k=0}^{n-1} \alpha_k \mathbb{1}_{[a+k\frac{1-a}{n}, a+(k+1)\frac{1-a}{n}]}$, $\alpha_i \geq 0$.
- $Z_i \triangleq (\beta_0^i, \dots, \beta_{d-1}^i) \in \mathbb{R}^d$, $\frac{1}{d} \sum_{k=0}^{d-1} \beta_k^i = 0$, $\frac{1}{d} \sum_{k=0}^{d-1} (\beta_k^i)^2 \leq 1$. Following the results presented in Appendix A, we have relaxed $\text{Var}[Z_i] = 1$ to $\text{Var}[Z_i] \leq 1$.

In this setting we have that

$$\varrho_i(W_i - Z_i \int_a^1 \sqrt{-v'(\theta)} f_i(\theta) d\theta) = r_i(\alpha, \beta, \gamma, \lambda).$$

and the problem amounts to minimizing the aggregate risk $\sum_{i=1}^2 r_i(\alpha, \beta, \gamma, \lambda)$ subject to $r_i(\alpha, \beta, \gamma, \lambda) \leq r_i(0, 0, 0, 0)$ and subject to constraints on α , β , γ and λ specified earlier. The latter are convex (in fact most of them are linear) in the corresponding variables. A crucial property of the functions $r_i : \mathbb{R}^{n+2d+1+n} \rightarrow \mathbb{R}$ is that for any three fixed entries, they are convex on the remaining variable. Exploiting the latter we use an iterative algorithm that performs a one-step, first-order descent in each direction alternately, such that the aggregate risk decreases in each step. The algorithm repeats this process until a maximal number of iterations has been performed or until there is no significant change in the aggregate risk. At each of the steps we encounter the problem of minimizing a convex function with respect to convex constraints.

Our coding has been done in Java, and we have used the NetBeans IDE 6.8 environment. In the case of *AV@R*, we have used the *or124.jar* package (OR-Objects 1.2.4) to calculate the required gradients. The descent procedures are all based on a local linearization of the function to be minimized, as well as of the constraints. Since all of these objects are convex, they can be locally, well approximated by




Figure 6: Dealing with non-feasible points

the corresponding subgradients (which are calculated using the chain rule). The linearized function is minimized, subject to the linearized constraints, on a cubic neighborhood of the current point. This reduces to a linear optimization problem. Again the linear optimization package included in *or124.jar* is used to obtain a minimizer. If this point does not satisfy the constraints (we carry the linearization error), a correction procedure is performed to obtain a feasible point close to the prior one. This procedure essentially performs a line search in the direction opposite to the direction of the subgradient of the constraint function at the point in hand (See Figure 6). Finally, it is verified whether at this feasible point the value of the objective function has decreased. Otherwise, the procedure is repeated starting with a cube whose size length is half of the original one.

C Convex risk measures on \mathbb{L}^2 .

In this appendix we recall some properties and representation results for risk measures on $\mathbb{L}^2(\Omega, \mathbb{P})$. We refer the reader to [7] and to section 4.3 in [15] for detailed discussions on this topic.

Definition C.1 (i) A monetary measure of risk on \mathbb{L}^2 is a function $\varrho : \mathbb{L}^2(\Omega, \mathbb{P}) \rightarrow \mathbb{R} \cup \{+\infty\}$ such that for all $X, Y \in \mathbb{L}^2$ the following conditions are satisfied:

- *Monotonicity:* if $X \leq Y$ then $\varrho(X) \geq \varrho(Y)$.
- *Cash Invariance:* if $m \in \mathbb{R}$ then $\varrho(X + m) = \varrho(X) - m$.

(ii) A risk measure is called coherent if it is convex and homogeneous of degree 1, i.e., if the following two conditions hold:

- *Convexity:* for all $\lambda \in [0, 1]$ and all positions $X, Y \in \mathbb{L}^2$:

$$\varrho(\lambda X + (1 - \lambda)Y) \leq \lambda\varrho(X) + (1 - \lambda)\varrho(Y)$$

- *Positive Homogeneity:* For all $\lambda \geq 1$

$$\varrho(\lambda X) = \lambda\varrho(X).$$

(iii) The risk measure is called law invariant, if

$$\varrho(X) = \varrho(Y)$$

for any two random variables X and Y which have the same law.

(iv) The risk measure ϱ on \mathbb{L}^2 has the Fatou property if for any sequence of random variables X_1, X_2, \dots that converges in norm to a random variable X we have

$$\varrho(X) \leq \liminf_{n \rightarrow \infty} \varrho(X_n).$$

Since a risk measure ϱ that has the Fatou property is a l.s.c. and proper convex mapping from $\mathbb{L}^2(\Omega, \mathbb{P})$ into $\mathbb{R} \cup \{+\infty\}$, we may represent it via a Legendre-Fenchel transform. Namely, if we define $\mathcal{M}_1(\mathbb{P})$ to be the set of probability measures on Ω that are absolutely continuous (w.r.t \mathbb{P}) and

$$\mathcal{M}_1^2(\mathbb{P}) \triangleq \left\{ Q \in \mathcal{M}_1(\mathbb{P}) \mid \frac{dQ}{d\mathbb{P}} \in \mathbb{L}^2 \right\},$$

then given ϱ there exists a penalty function $\alpha : \mathbb{L}^2(\Omega, \mathbb{P}) \rightarrow \mathbb{R} \cup \{+\infty\}$ such that

$$\varrho(X) = \sup_{Q \in \mathcal{M}_1^2(\mathbb{P})} \{ \mathbb{E}_{\mathbb{P}}[-X] - \alpha(Q) \}.$$

Moreover, if for $Q \in \mathcal{M}_1^2(\mathbb{P})$ we write $Y = dQ/d\mathbb{P}$, then the above expression can be written as

$$\varrho(X) = \sup_{\|Y\|_2=1} \{ -\langle X, Y \rangle - \beta(Y) \},$$

where $\langle \cdot, \cdot \rangle$ is the canonical inner product in $\mathbb{L}^2(\Omega, \mathbb{P})$, and $\beta(Y) = \alpha(Q)$.

References

- [1] C. ALIPRANTIS & K. BORDER: Infinite Dimensional Analysis, a Hitchhiker's guide (3rd. edition), *Springer Verlag*, 2006.
- [2] M. ARMSTRONG: Multiproduct Nonlinear Pricing, *Econometrica*, 64, 51-75, 1996.
- [3] P. ARTZNER, F. DELBAEN, J-M EBERT & D. HEATH: Coherent Measures of Risk, *Mathematical Finance*, vol. 9, no. 3, 203-228, 1999.
- [4] A. BAGH & A. JOFRÉ: Reciprocal Upper Semicontinuity and Better Reply Secure Games, a Comment, *Econometrica* vol. 74, no. 6, pp. 1715-1721, 2006.
- [5] P. BARELLI & I. SOZA: On the Existence of Nash Equilibria in Discontinuous and Qualitative Games, mimeo.

- [6] P. BARRIEU & N. EL KAROUI: Optimal risk transfer, *Finance*, 25, pp. 31-47, 2004.
- [7] P. CHERIDITO & T. LI: Dual Characterization of Properties of Risk Measures in Orlicz Hearts, *Mathematics and Financial Economics*, vol. 2, no. 1, pp. 29-55, 2008.
- [8] P. CHERIDITO, U. HORST, M. KUPPER & T.A. PIRVU: Equilibrium Pricing in Incomplete Markets under Translation Invariant Preferences, *Working Paper*.
- [9] G. CARLIER, I. EKELAND & N. TOUZI: Optimal Derivatives Design for Mean-Variance Agents under Adverse Selection, *Mathematics and Financial Economics*, vol. 1, no. 1, pp. 57-80, 2007.
- [10] P. DASGUPTA & E. MASKIN: The Existence of Equilibrium in Discontinuous Economic Games, I: Theory, *Review of Economic Studies*, vol. 53, no. 1, pp. 1-26, 1986.
- [11] I. EKELAND & S. MORENO-BROMBERG: An Algorithm for Computing Solutions of Variational Problems with Global Convexity Constraints, *Numerische Mathematik*, 2009.
- [12] I. EKELAND & R. TÉMAM: Convex Analysis and Variational Problems, *Classics in Applied Mathematics*, 28, SIAM, 1976.
- [13] L. EPSTEIN & M. PETERS: A Revelation Principle for Competing Mechanisms, *Journal of Economic Theory*, 88, pp. 119-160, 1999.
- [14] D. FILIPOVIĆ & M. KUPPER: Equilibrium Prices for Monetary Utility Functions, *International Journal of Theoretical and Applied Finance*, 11, pp. 325 - 343, 2008.
- [15] H. FÖLLMER & A. SCHIED: Stochastic Finance. An Introduction in Discrete Time, *de Gruyter Studies in Mathematics*, 27, 2004.
- [16] H. FÖLLMER & A. SCHIED: Convex Measures of Risk and Trading Constraints, *Finance and Stochastics*, vol. 6, no. 4, pp. 429-447, 2002.
- [17] M. FRITTELLI & E. ROSAZZA GIANIN: Putting Order in Risk Measures, *Journal of Banking & Finance*, 26, pp. 1473-1486, 2002.
- [18] D. FUDENBERG & J. TIROLE: Game Theory, *MIT Press*, 1991.
- [19] R. GUESNERIE: A contribution to the Pure Theory of Taxation, *Econometrica*, 49, pp. 33-64, 1995.
- [20] U. HORST & S. MORENO-BROMBERG: Risk Minimization and Optimal Derivative Design in a Principal Agent Game, *Mathematics and Financial Economics*, vol. 2, no. 1, pp. 1-27, 2008.
- [21] E. JOUINI, W. SCHACHERMEYER & N. TOUZI: Optimal Risk Sharing for Law Invariant Monetary Utility Functions, *Mathematical Finance*, vol. 18, no. 2, pp. 269-292, 2008.
- [22] P. MILGROM & I. SEGAL: Envelope Theorems for Arbitrary Choice Sets, *Econometrica*, vol. 70, no. 2, pp. 583-601, 2002.
- [23] M. MUSSA & S. ROSEN: Monopoly and Product Quality, *Journal of Economic Theory*, 18, pp. 301-317, 1978.
- [24] D. MARTIMORT & L. STOLE: The Revelation and Delegation Principles in Common Agency Games, *Econometrica*, vol. 70, no. 4, pp. 1659-1673, 2002.

- [25] R. NESSAH & G. TIAN: The Existence of Equilibria in Discontinuous and Nonconvex Games, working paper.
- [26] F. PAGE JR.: Catalogue Competition and Stable Nonlinear Prices, *Journal of Mathematical Economics*, 44, pp. 822-835, 2008.
- [27] F. PAGE JR. & P. MONTEIRO: Three Principles of Competitive Nonlinear Pricing, *Journal of Mathematical Economics*, 39, Issues 1-2, pp. 63-109, 2003.
- [28] F. PAGE JR. & P. MONTEIRO: Uniform Payoff Security and Nash Equilibrium in Compact Games, *Journal of Economic Theory*, (to appear)
- [29] F. PAGE JR. & P. MONTEIRO: Catalogue Competition and Nash Equilibrium in Monlinear Pricing Games, *Economic Theory*, 34, pp. 503-524, 2008.
- [30] R. PHELPS: Convex Functions, Monotone Operators and Differentiability:, *Lecture Notes in Mathematics*, 1364, Springer-Verlag, 1989.
- [31] P. RENY: On the Existence of Pure and Mixed Strategy Equilibria in Discontinuous Games, *Econometrica*, 67, pp. 1029-1056, 1999.
- [32] J.-C. ROCHET: A Necessary and Sufficient Condition for Rationalizability in a Quasi-linear Context, *Jornal of Mathematical Economics*, 16, pp. 191-200, 1987.
- [33] J.-C. ROCHET & P. CHONÉ: Ironing, Sweeping and Multidimensional Screening, *Econometrica*, 66, pp. 783-826, 1988.
- [34] G. TIAN: The Existence of Equilibria in Games with Arbitrary Strategy Spaces and Payoffs: A Full Characterization, working paper.