

(Neutrally) Optimal Mechanism under Adverse Selection: the canonical insurance problem

Theodoros M. Diasakos*

Kostas Koufopoulos[‡]

May 26, 2018

Abstract

This paper revisits the problem of adverse selection in the insurance market of Rothschild and Stiglitz [35]. We extend the three-stage game in Hellwig [19] by allowing firms to endogenously choose whether or not to pre-commit on their contractual offers (menus). We show how this mechanism can deliver the Miyazaki-Wilson-Spence allocation as the unique perfect-Bayesian equilibrium. This allocation is the unique incentive-efficient and individually-rational maximizer of the utility of the most profitable type. In fact, given that the informed player has only two types, it is the unique core allocation and thus the unique neutral optimum in the sense of Myerson [29].

Keywords: Insurance Market, Adverse Selection, Interim Incentive Efficiency, Neutral Optimum.

JEL Classification Numbers: D86

Accepted refereed manuscript of:

Diasakos TM & Koufopoulos K (2018) (Neutrally) Optimal Mechanism under Adverse Selection: The canonical insurance problem, *Games and Economic Behavior*, 111, pp. 159-186.

DOI: [10.1016/j.geb.2018.04.007](https://doi.org/10.1016/j.geb.2018.04.007)

© 2018, Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>

*Economics Division, Stirling Management School, University of Stirling

[†]University of Leicester School of Business.

[‡]We have benefited greatly from comments and suggestions by the associate editor and two anonymous referees as well as by Aris Boukouras, Tasos Dosis, Dino Gerardi, Dominik Grafenhofer, Peter Hammond, Martin Hellwig, Miltos Makris, Eric Maskin, Motty Perry, Heracles Polemarchakis, Philip Reny, Paolo Siconolfi, and Yiannis Vailakis. Diasakos acknowledges also financial support by the Scottish Institute for Research in Economics. Any errors are ours.

1 Introduction

An old yet still open issue in applied micro-economic theory concerns the nature of competitive equilibrium in insurance markets under adverse selection. This dates back at least to the seminal study in Rothschild and Stiglitz [35], and has led to a significant body of research which exhibits renewed momentum in the wake of the last financial crisis. Restricting attention to the insurance market in Rothschild and Stiglitz [35] - what we will be referring to henceforth as the “canonical insurance problem” - the present paper proposes a game-theoretic mechanism in which the conjunction of Bertrand competition, on the one hand, and Nash-type strategic interaction, on the other, delivers the strongest of results.

Namely, the Miyazaki-Wilson-Spence (MWS) allocation emerges as the unique perfect Bayesian equilibrium (PBE) outcome. As is well-known, this allocation entails the separation of risk-types across insurance contracts, may involve cross-subsidization, and is always interim incentive efficient (IIE). Being in fact the most preferred point on the IIE frontier for the profitable risk-type, it is the most desirable allocation with respect to core and neutral optimality (either notion in the sense of Myerson [29] and to be described in the sequel).

More specifically, we consider the following extension of the three-stage game in Hellwig [19]. At stage 1, the firms (insurers) simultaneously offer collections of insurance menus, specifying for each menu in their collection whether or not they commit upon its delivery. At stage 2, having observed which menus have been offered with commitment, the customers (insurees) choose amongst the firms’ offers. At stage 3, after observing the offers of stage 1 (including whether or not a given menu has been offered with commitment) but also whether or not a given menu has been chosen, the firms that offered menus without commitment decide simultaneously whether or not to withdraw them. We show that this mechanism can deliver the MWS menu (offered without commitment) as the unique equilibrium allocation.

The key feature of our mechanism is *endogenous commitment*: firms can choose to *publicly commit* at stage 1 to deliver a given menu (allocation) at stage 3, irrespectively of the history of play at that point. Intuitively, the fact that a firm can publicly commit to deliver a given menu means that the customers’ decision of whether or not to choose it depends only on preferences, not on beliefs - for they need not worry about the possibility that the menu might be later withdrawn (an action that would depend on the firm’s belief about the distribution of risk-types in its pool of customers, which would in turn depend on the actions and beliefs of the other firms and the customers, etc.). As a result, starting from any allocation that is not the MWS one but satisfies the constraints for IIE optimality, it is always possible for a firm to offer with commitment a cream-skimming menu that remains profitable even if it attracts both risk-types (see Lemma A.3). That is, the combination of the firms’ ability to cross-subsidize and to commit publicly to a menu (off the equilibrium path) leaves the MWS as the unique candidate equilibrium allocation (see Claim 3.1). As there are no allocations that can cream-skin against the MWS menu and remain profitable when attracting both risk-types, the threat of withdrawing the MWS menu fends off cream-skimming deviations with commitment. Regarding cream-skimming deviations without commitment, the

signalling nature of the subgame in stages 2-3 supports an equilibrium scenario in which the MWS menu is never withdrawn at stage 3 (see Proposition 3.1).

It is noteworthy that the market dynamics described above lie within the realm of Bertrand competition, embedding in fact the seminal studies in Rothschild and Stiglitz [35] and Hellwig [19]. The two models result directly from our mechanism if we restrict the firms to making offers at stage 1, respectively, only with and only without commitment. That is, the models in Rothschild and Stiglitz [35] and Hellwig [19] can be viewed, respectively, as examples of Bertrand competition *without exit* and *with exit but without the ability to publicly commit to not exit*. As we show in the sequel, this perspective sheds light in important aspects of Bertrand competition under adverse selection that seem to have been ignored by the two papers, causing some confusion in the subsequent literature. Contrary to what seems to be a common belief in the literature, the restriction either model imposed on competition under adverse selection went far beyond not allowing firms to cross-subsidize between risk-types.

If firms can make offers only with commitment, the impossibility of exit forces them to consider only menus that are *safe* in the sense of Holmstrom and Myerson [20]: in the current context, separating menus that avoid losses on a contract-by-contract (i.e., type-by-type) basis; in fact, on the equilibrium path they must break even on a contract-by-contract basis (see Lemma B.2). Moreover, the conjunction of free-entry with the fact that firms can offer cross-subsidizing menus only with commitment renders the MWS the unique candidate equilibrium allocation. Clearly, if firms can cross-subsidize but not exit, the MWS menu is the unique equilibrium allocation when it coincides though with the Rothschild-Stiglitz-Wilson (RSW) allocation (the unique candidate equilibrium in Rothschild and Stiglitz [35]); equivalently, only if the RSW allocation is IIE (see Claim 3.1 in the Online Appendix). Yet the impossibility of withdrawal renders the RSW allocation vulnerable to cream-skimming deviations, leading to the issue of non-existence of equilibrium in Rothschild and Stiglitz [35]. As it turns out, there are no profitable cream-skimming deviations if the RSW allocation is IIE (see Claim 4.2).¹

If, on the other hand, the firms can make offers only without commitment, any allocation that is individually rational for the firms can be supported now as equilibrium (see Claim 4.1). The multiplicity of equilibria arises because deviant menus can be offered only without commitment: in the signalling subgame of stages 2-3, the possibility that the deviant menu is withdrawn renders its expected profitability a function of the players' beliefs about its profitability - and beliefs can be arbitrary off the equilibrium path. Clearly, the real issue for Bertrand competition under adverse selection in Rothschild and Stiglitz [35] has to do solely with the existence of equilibrium. By contrast, the real issue in Hellwig [19] has to do solely with the efficiency of the equilibrium allocation. By combining free entry and exit with cross-subsidization and endogenous commitment, our approach addresses either issue simultaneously. And suggests that the real desideratum for market competition under adverse selection ought to be enabling the firms not only to freely enter and exit, but also to publicly commit to not exit.

¹In Rothschild and Stiglitz [35], the fact that the space of menus (and thus also that of deviations) is further restricted to not allow cross-subsidization relaxes slightly the issue of non-existence of equilibrium (see Claim 4.3).

In terms of the fundamentals of cross-subsidization, our approach builds upon the insights in Miyazaki [28] and Grossman [17]. The former paper was the first to introduce menus of contracts as the objects of trade: to allow that is firms to offer menus of contracts but withdraw a given contract only if they do so for the entire respective menu. The main idea in fact was to identify a given firm by the menu it offers: menu withdrawals are interpreted as the corresponding firms exiting the market entirely. By contrast, the latter paper suggested that, especially within the context of insurance provision, one ought to imagine customers submitting applications and firms deciding whether or not to approve the applications they receive.

Taking this as our interpretation of the last two stages in Hellwig’s game, we introduce two additional elements in the applications’ process. Our firms may offer (finite) collections of menus, and subsequently withdraw single menus from their collections. They may also choose to send their customers “pre-approved” application forms. The former element means that “exit” may refer to a given menu from a firm’s collection (a particular product from a firm’s line of products), not necessarily to the firm closing down. The second element is an equally realistic insight on common practice in the markets for standardized insurance or credit provision. It readily admits moreover a game-theoretic interpretation, as a *public action* in the terminology of Myerson [29]: an enforceable decision a player can publicly-commit to carry out, even if it may turn out ex-post to be harmful to her or others.²

In terms of related studies, the underlying issue being so fundamental there is a large and important literature on competitive equilibria under adverse selection. Its principal aim has been to propose mechanisms, along with their implementing market structures, to ensure that always some allocation will be supported as competitive equilibrium, under some associated notion of competitive equilibrium. The respective models can be broadly classified into three categories, based upon the extent to which the mechanism allows the players’ behavior to be strategic.

One class of models (see, amongst others, Prescott and Townsend [32]-[33], Rustichini and Siconolfi [36], Dubey and Geanakoplos [12], Dubey et al. [13], and Guerrieri et al. [18]) has focused on Walrasian mechanisms. The central message has been that general economies with adverse selection do not always admit pure Walrasian equilibrium pricing systems, and when they do the resulting allocations are not necessarily IIE. To guarantee existence of equilibrium some studies have introduced rationing or suppressed the requirement that firms are profit-maximizing while imposing at the same time quantity constraints on trade.³ Either of these approaches arrives at

²Many a reader will have at some point received in the mail application forms, especially from credit-card companies, that name the recipient “pre-approved” for a particular product. Many a reader will have also filled in and mailed back such a form, only to find out that the small print one typically misses renders the term “pre-approved” devoid of any meaning other than “pre-selected as a marketing target.” In sharp contrast, our mechanism demands explicitly that the term is legally binding. And prescribes implicitly the required supporting institutional substructure as a sufficient condition for efficient insurance provision in the presence of adverse selection.

³A notable exception is Bisin and Gottardi [8] where, instead of being constrained, the Walrasian mechanism is enhanced with the implicit presence of institutions that monitor trade appropriately. Restricting attention also to the canonical insurance economy, this paper shows that the RSW allocation obtains always as the unique Walrasian equilibrium if there are markets for contingent claims in which agents trade only incentive-compatible contracts. Yet,

some equilibrium that is essentially unique and involves a separating allocation. Typically, however, this is not IIE while uniqueness obtains by restricting the out-of-equilibrium actions and beliefs often in strong ways.⁴

Another perspective has been to look at mechanisms in which competitive equilibrium is supported by strategic behavior, and this has produced two separate lines of approach (for a short yet comprehensive review, see Mimra and Wambach [26]). In the early models (see Wilson [40], Miyazaki [28], Grossman [17], Riley [34], but also Engers and Fernandez [14]), some of the players exhibit strategic behavior which is not of the Nash-type.⁵ Because of this, a most important shortcoming of this strand of the literature is the absence of game-theoretic structure.⁶ In sharp contrast, more recent studies (see Netzer and Scheuer [30], Picard [31], von Siemens and Kosfeld [37]-[38], as well as Asheim and Nilssen [3]) rest upon the game-theoretic foundation in Hellwig [19] (or its generalization in Maskin and Tirole [25]). As such, they are directly related to the present paper and will be discussed in some detail in the sequel.

The balance of the paper is structured as follows. The next section introduces the canonical insurance problem and recalls some useful results concerning the efficiency of the feasible allocations. Section 3 describes the mechanism and explains intuitively its equilibrium dynamics. In Section 4, we interpret our mechanism through the prism of the directly-related results in the literature - demonstrating in particular how our approach sheds light to the issues about competition under adverse selection discussed above. Section 5 concludes, being followed by appendices that contain the proofs of the results in the main text (supporting technical material and results are to be found in the Online Appendix).

to ensure that incentive efficiency is attained even when the RSW allocation is not IIE, markets for consumption rights are required. The ensuing Arrow-Lindahl equilibria internalize the consumption externality due to adverse selection. By varying the endowment of consumption rights, the equilibrium set traces the entire IIE frontier but for the MWS point. The latter can be obtained only at the limit of a sequence of economies as the distribution of consumption rights varies accordingly.

⁴In its pure form, the Walrasian approach was initiated by Prescott and Townsend [32]-[33] and revisited recently by Rustichini and Siconolfi [36]. Under rationing, the refinement criteria range from subgame perfection (see Guerrieri et al. [18]) to the universal divinity of Banks and Sobel [6] (see Gale [15]) or the stability of Kohlberg and Mertens [22] (see Gale [16]). The latter notion has been deployed also under quantity constraints (see Dubey and Geanakoplos [12] or Dubey et al. [13]) but seems to be more binding in that environment - as shown in Martin [23], its weakening to something akin to trembling-hand perfection leads to many additional equilibria.

⁵The sellers in Wilson [40], Miyazaki [28], Riley [34], Engers and Fernandez [14] but also the buyers in Grossman [17] are able to foresee the unraveling of equilibrium Rothschild and Stiglitz [35] warned about and modify their plans so as to prevent it.

⁶A notable exception here is Engers and Fernandez [14]. Allowing for cross-subsidization, this study extends the notion of reactive equilibrium in Riley [34] to one that admits the RSW allocation as the unique outcome. The authors present also a game, with an infinite number of moves, in which the RSW allocation is supported as SPNE. Alas, so are a very large collection of other allocations within which not only the RSW allocation but also the IIE frontier are subsets of negligible relative size.

2 The Canonical Insurance Problem

Consider a continuum (of unit measure) of risk averse agents, each having the same wealth endowment $W \in \mathbb{R}_{++}$ and facing the idiosyncratic risk of it being reduced by the amount $d \in (0, W)$. The endowment risk is depicted by two states of nature, “no accident” and “accident” (denoted, respectively, by $s \in \{0, 1\}$), with either state perfectly observable and verifiable while each agent belongs to one of two risk classes, “low-risk” and “high-risk” (denoted, respectively, by $h \in \{L, H\}$). The class of an agent (her “risk-type”) corresponds to the probability of having an accident (with $0 < p_L < p_H < 1$). It is her own private information whereas the share $\lambda \in (0, 1)$ of low-risk agents in the population is common knowledge.⁷

An agent may insure herself against her endowment risk by accepting an insurance contract. The space of admissible contracts is $\mathcal{A} = [-W_0, W_1] \times [-W_0, W_1]$, for some constants $(W_0, W_1) \in \mathbb{R}_+ \times \mathbb{R}_{++}$. The typical contract $\mathbf{a} = (a_0, a_1) \in \mathcal{A}$ determines the premium a_0 to be paid in exchange for receiving the net indemnity a_1 if an accident occurs. Letting the function $\mathbf{w}(\mathbf{a}) = (w_0(a_0), w_1(a_1)) := (W - a_0, W - d + a_1)$ depict the corresponding wealth vector gives a one-to-one relation between \mathcal{A} and the wealth space $\mathcal{W} = [W - W_1, W + W_0] \times [W - d - W_0, W - d + W_1]$. It entails thus only slight abuse of notation to take the agents’ preference relation over lotteries on \mathcal{W} , denoted by \succsim_h for risk-type h , as ordering also lotteries on \mathcal{A} .

Regarding these preferences, partitioning the population into the low- and high-risk classes is supposed to capture all the relevant information about the problem of adverse selection (i.e., the partition obtains after having inferred as much as possible from the agents’ observable characteristics). As a result, apart from being of a particular risk-type, the agents are otherwise identical. They all have von Neumann-Morgenstern preferences over wealth lotteries, the expected utility of an agent of risk-type h being given by $U_h(\mathbf{w}(\mathbf{a})) = (1 - p_h)u(w_0(a_0)) + p_h u(w_1(a_1))$. The state- and type-independent Bernoulli utility function $u : [W - d - \bar{W}, W + \bar{W}] \rightarrow \mathbb{R}$, where $\bar{W} = \max\{W_0, W_1 - d\}$, will be taken to be everywhere strictly-increasing, strictly-concave, and twice continuously-differentiable. And to abuse notation slightly further, we will often take $U_h(\cdot)$ as being defined directly on \mathcal{A} .

On the supply side of the market, insurance is provided by risk-neutral firms which maximize expected profits. For the typical firm this is given by $\Pi_h(\mathbf{a}) = (1 - p_h)a_0 - p_h a_1$ when the contract is bought by a customer of risk-type h . To ensure competition, we will assume that the collection of firms forms a finite set \mathcal{N} with cardinality $|\mathcal{N}| \geq 2$. We will suppose also that the insurance market in question admits free entry: the firms have adequate financial resources to be able to deliver, following the realization of the accident state, any collection of contracts from \mathcal{A} to any subset of the population.

⁷The continuum hypothesis is standard in models of this kind. It facilitates the deployment of some version of the law of large numbers to enable the assumption that exactly the shares p_L [resp. p_H] of the low-risk [resp. high-risk] agents and $\lambda p_L + (1 - \lambda)p_H$ of the entire population will eventually have an accident. Of course, applying laws of large numbers on a continuum of random variables is not without technical caveats - see the footnote 3 in Netzer and Scheuer [30] for some references on the technical complications and the respective remedies.

Feasibility and Efficiency

Given this data, we can determine the collection of contracts that are relevant for the problem at hand under the standard assumption that the choice of contract by an agent is *exclusive*.⁸ Each risk-type, that is, will be restricted to choose at most one contract from those available in the market when she is called upon to act. Under this condition, since the population is partitioned into two risk classes, the relevant collection of contracts is the space of douplets $\{\mathbf{a}_L, \mathbf{a}_H\}$ (which we will call insurance *menus*) where $\mathbf{a}_L, \mathbf{a}_H \in \mathcal{A}$ while the subscript indicates the risk-class the respective contract is meant for. To ensure though that we remain within the realm of the direct revelation principle, and that market participants are rational with an outside option available, we will restrict the relevant space of insurance menus to those that are *incentive compatible* for either risk-type and *individually rational* for all market participants.

The typical menu $m = \{\mathbf{a}_L, \mathbf{a}_H\}$ meets the former requirement if it satisfies⁹

$$(\text{IC}_h) : \quad U_h(\mathbf{a}_h) - U_h(\mathbf{a}_{h'}) \geq 0, \quad h, h' \in \{L, H\}, h' \neq h$$

while it is individually-rational for risk-type h if it satisfies¹⁰

$$(\text{IR}_h) : \quad U_h(\mathbf{a}_h) \geq \bar{u}_h \equiv U_h(\mathbf{0}), \quad h \in \{L, H\}$$

As for the firms, they can guarantee themselves no losses simply by abstaining from market activity. The menu is thus individually-rational for a firm if it avoids losses at least for some belief the firm may have about the distribution of risk-types in its customer pool:

$$(\text{IR}_F) : \quad \exists p \in [0, 1] : \Pi(m|p) := p\Pi_L(\mathbf{a}_L) + (1-p)\Pi_H(\mathbf{a}_H) \geq 0$$

The space of relevant menus is thus given by

$$\mathcal{M} = \{ \{ \mathbf{a}_L, \mathbf{a}_H \} : \mathbf{a}_L, \mathbf{a}_H \in \mathcal{A}, (\text{IC}_L), (\text{IC}_H), (\text{IR}_L), (\text{IR}_H), (\text{IR}_F) \}$$

⁸Even though widely used in the literature, this is admittedly a strong requirement. Non-exclusive contracting in the context of an insurance market under adverse selection has been considered in the literature only very recently. With respect to the canonical insurance problem, Ales and Maziero [1] introduce a third risk-class of customers (the medium-risk type), and conduct their study using free-entry arguments. Attar et al. [4] on the other hand allow for more general risk-preferences for either risk-type, and proceed with a game-theoretic analysis. The differences between the two models notwithstanding, the common key insight is that a pure-strategy equilibrium may fail to exist; and when it does exist, insurance is provided only to the high-risk customers (in Attar et al. [4], it may also require that the low-risk type becomes herself an insurer).

⁹This condition ensures that truthfully-revealing one's risk-type is a Nash equilibrium amongst the customers: it is in one's best interest to report her type honestly given that everyone else does the same. A menu that is incentive compatible for either risk-type can be either *separating* (if at least one of the (IC_h) inequalities is strict) or *pooling* (if either of the (IC_h) constraints binds). The latter type of menus cannot but be of the form $\{\mathbf{a}, \mathbf{a}\}$ for some $\mathbf{a} \in \mathcal{A}$ (see Lemma 1.8 in the Online Appendix).

¹⁰This condition reflects the very fact that each customer enters the insurance market while still in possession of the endowment lottery. Technically, it suffices to assume a zeroth firm with a single plan of action: offering and honoring just the menu $\{\mathbf{0}, \mathbf{0}\}$ irrespectively of the history of play. Observe also that we may let $\bar{u}_h > 0$ with no loss of generality: since $U_h(\mathbf{0}) = (1-p_h)u(W) + p_h u(W-d)$, normalizing the customers' outside option in the accident state to $u(W-d) = 0$ gives $U_h(\mathbf{0}) = (1-p_h)u(W) > (1-p_h)u(W-d) = 0$.

and within it the search for optimal menus will be based of course upon the Pareto criterion. Given two menus $m = \{\mathbf{a}_L, \mathbf{a}_H\}$ and $m' = \{\mathbf{a}'_L, \mathbf{a}'_H\}$, the former Pareto-dominates the latter weakly if $U_h(\mathbf{a}_h) \geq U_h(\mathbf{a}'_h)$ for either h , and strictly if at least one inequality is strict. As is well-known, moreover, it suffices here to search for the *incentive efficient* (IE) allocations. These are the Pareto-optimal amongst the allocations that satisfy (IC_h) for either h and (IR_F) for $p = \lambda$. Formally, the IE frontier consists of those menus that solve the $IE(\mu)$ -problem:¹¹

$$\max_{\{\mathbf{a}_L, \mathbf{a}_H\}=m \in \mathcal{M}_\lambda} \mu U_L(\mathbf{a}_L) + (1 - \mu) U_H(\mathbf{a}_H)$$

where $\mu \in [0, 1]$ and

$$\mathcal{M}_\lambda = \{m \in \mathcal{M} : \Pi(m|\lambda) \geq 0\}$$

In what follows, we will be referring to an allocation as IE if it is the $IE(\mu)$ -optimum for some $\mu \in [0, 1]$. The respective menu will be denoted as $m^\mu = \{\mathbf{a}_L^\mu, \mathbf{a}_H^\mu\}$.

Our main focus though will be on two allocations that feature prominently in the literature. The Rothschild-Stiglitz-Wilson (RSW) allocation is the only allocation that solves the $RSW(\mu)$ -problem:

$$\max_{\{\mathbf{a}_L, \mathbf{a}_H\}=m \in \mathcal{M}_\lambda} \mu U_L(\mathbf{a}_L) + (1 - \mu) U_H(\mathbf{a}_H) : \quad \Pi_h(\mathbf{a}_h) \geq 0, \quad h \in \{L, H\}$$

for *every* $\mu \in [0, 1]$. This maximizes the utility of either risk-type while being also *ex-post* individually-rational for the firms. Its importance lies in the latter characteristic, which renders it a *safe* allocation in the terminology of Myerson [29]: one the uninformed party would continue to view as individually-rational even if she knew the private information of the informed, whatever this information might be. As is well-known, in the current setting the RSW allocation is given by the separating menu $m^{**} = \{\mathbf{a}_L^{**}, \mathbf{a}_H^{**}\}$, which offers to the high-risk type her full-insurance first-best allocation: $\mathbf{a}_H^{**} = \arg \max_{\mathbf{a} \in \mathbf{A} : \Pi_H(\mathbf{a}) \geq 0} U_H(\mathbf{a})$ (which is also such that $a_{0H}^{**} = d - a_{1H}^{**}$, $U_H(\mathbf{a}_L^{**}) = U_H(\mathbf{a}_H^{**})$, and $\Pi_h(\mathbf{a}_h^{**}) = 0$ for $h \in \{L, H\}$).¹²

¹¹The appropriateness of the IE criterion stems from that there cannot be unanimous agreement to depart from an IE allocation if individuals know just their own private information (see Theorem 1 and the subsequent discussion in Holmstrom and Myerson [20]). Notice also that we may represent the Pareto-frontier via changing the parameter μ in the objective of the $IE(\mu)$ problem as long as the feasible set is convex (see Proposition 16.E.2 in Mas-Colell et al. [24]). And even though the profit constraint is trivially convex, the incentive-compatibility ones are not. It is thus standard practice in the literature (see for instance Myerson [29] or Maskin and Tirole [25]) to assume that allocations are assigned randomly. In the present setting, this refers to distributions on $\mathcal{A} \times \mathcal{A}$ and, letting v be a typical one, makes the (IC_h) constraint read $\int_{\mathcal{A} \times \mathcal{A}} [U_h(\mathbf{a}_h) - U_h(\mathbf{a}_{h'})] dv(\mathbf{a}_h, \mathbf{a}_{h'}) \geq 0$; being linear, this is trivially convex in v . As it turns out, the IE frontier consists here entirely of deterministic allocations so that this assumption can be left in the background as a purely technical caveat.

¹²In the terminology of Maskin and Tirole [25], this is the RSW allocation relative to the endowment while the IE optima are interim efficient (and, thus, also weakly interim efficient) relative to the prior belief $\{\lambda, 1 - \lambda\}$. In the present study, however, the prior and interim beliefs will always coincide: the uninformed party will always be the one to make the proposals and thus no updating of beliefs will ever take place at the interim stage. This is why we choose to ignore the term “interim” all together.

Even more of interest in the sequel will be what is known as the Miyazaki-Wilson-Spence (MWS) allocation, which solves the IE(1)-problem but for an additional restriction on the feasible set. The restriction strengthens the individual-rationality of the high-risk type (whose presence is the very cause of the adverse selection problem) by requiring that her participation in the insurance market renders her at least as well off as if she were to reveal her type. Formally, letting her individual-rationality constraint be now given by

$$(\text{IR}_H^{\text{FB}}) : \quad U_H(\mathbf{a}_H) \geq U_H(\mathbf{a}_H^{**})$$

the MWS allocation solves what we will be referring to as the MWS problem:

$$\max_{\{\mathbf{a}_L, \mathbf{a}_H\} = m \in \mathcal{M}_\lambda} U_L(\mathbf{a}_L) : \quad (\text{IR}_H^{\text{FB}})$$

The MWS allocation is unique (see Lemma 2.1 in the Online Appendix), given by the IE($\bar{\mu}$)-optimum where $\bar{\mu} = \max \{\mu \in (\lambda, 1] : U_H(\mathbf{a}_H^\mu) \geq U_H(\mathbf{a}_H^{**})\}$ (see Lemma 2.6 in the Online Appendix). It will be denoted henceforth by $\bar{m} = \{\bar{\mathbf{a}}_L, \bar{\mathbf{a}}_H\}$.

Remark 1. The importance of the MWS menu emerges when we combine the following facts: (i) either type's utility varies monotonically with μ on the IE frontier (see Lemma 1.5 in the Online Appendix), and (ii) the low-risk [resp. high-risk] type is the profitable type on $[0, \bar{\mu}]$ [resp. on $[0, 1] \setminus [0, \bar{\mu}]$] (see Lemma 2.4 in the Online Appendix). The MWS allocation is the point on the IE frontier that maximizes the utility of the profitable type.

Remark 2. Investigating the issue of allocative efficiency, the original definition of the IE(μ)-problem does not actually require that the (IR_h) constraints are met. When it comes though to the issue of efficient insurance provision in a competitive market under free entry, these become relevant conditions for the customers' participation. Nevertheless, and in sharp contrast to most of the recent literature on the canonical insurance problem (save for the models that result in equilibrium sets consisting only of the IE allocations that Pareto-dominate the RSW one), this discrepancy will be of no consequence for the subsequent analysis. Both in terms of supporting the equilibrium scenario as well as of establishing its uniqueness, our arguments will deploy menus that lie arbitrarily close to the MWS one. And such menus are guaranteed to satisfy either (IR_h) constraint (see Lemma 2.5 and Remark IV in the Online Appendix). They lie in fact in the space

$$\mathcal{M}_\lambda^+ = \left\{ \{\mathbf{a}_L, \mathbf{a}_H\} = m \in \mathcal{M} : \Pi(m|\lambda) > 0, \mathbf{a}_{h \in \{L, H\}} \in \mathbb{R}_{++}^2 \right\}$$

3 The Mechanism

Our mechanism operates as follows. At stage 1, the firms decide simultaneously on their offers. Each firm must choose a finite collection from \mathcal{M} and, at the same time, whether or not to pre-commit upon the delivery of each menu in this collection.

At stage 2, from the total set of contractual offers the firms have introduced at stage 1, the customers choose simultaneously menus to apply for and contracts to receive out of these menus.

Their decisions are made after observing whether or not a given menu has been offered with commitment. Each risk-type can apply for up to one menu, and choose only one contract from that menu. If a customer doesn't apply for any menu, she is left to face her endowment lottery. If no menu receives applications, the game ends. If a customer applies for a menu that has been offered with commitment, she receives her chosen contract from that menu. If all the menus introduced at stage 1 have been offered with commitment or all customers at stage 2 apply for menus offered with commitment, the game ends; otherwise, it continues to the next stage.

At stage 3, after observing the offers of stage 1 - including whether or not a given menu has been offered with commitment, - the firms that offered menus without commitment decide simultaneously whether or not to withdraw them. They are allowed to withdraw single menus from their collection, but not single contracts from menus. The firms make their decisions after observing whether or not a given contract in a given menu has received applications. If the menu a customer had applied for at stage 2 gets withdrawn at stage 3, she is left to face her endowment lottery; otherwise, she receives her chosen contract from that menu. In either case, the game ends.

We will denote the above game by $\Gamma = (S_F, S_A, \Pi_{h \in \{L, H\}}, U_{h \in \{L, H\}})$. Pure-strategy profiles for the firms will be denoted by $s = (s_1, \dots, s_N) \in S_F$ with the typical strategy having three components. The first two reflect the action of the typical firm $n \in \mathcal{N}$ at stage 1: the choice $\mathcal{M}_n \in \overline{\mathcal{M}} := \{M \in 2^{\mathcal{M}} : |M| \in \mathbb{N}\}$ and, for each $m \in \mathcal{M}_n$, the decision of whether or not to pre-commit upon its delivery. This will be depicted by the function $c : \mathcal{M}_n \rightarrow \mathcal{C} \equiv \{C, NC\}$. We will let $\cup_{m \in \mathcal{M}_n} (m, c(m)) \equiv s_n^1 \in S_n^1 \equiv \cup_{M \in \overline{\mathcal{M}}} M \times \mathcal{C}$. Abusing notation slightly, we will let also

$$s_{n|\mathcal{M}}^1 = \mathcal{M}_n \quad s_{n|C}^1 = \left\{ m \in s_{n|\mathcal{M}}^1 : c(m) = C \right\}$$

The typical history to be observed at the beginning of stage 2 will be $(s_1^1, \dots, s_N^1) = s^1 \in S^1 = \prod_{n \in \mathcal{N}} S_n^1$. This determines $\Gamma(s^1)$, the (signalling) subgame of stages 2-3.¹³ The third component of the typical firm's strategy is given by a functional $s_n^2 : S^1 \rightarrow \{W, NW\}^{\overline{\mathcal{M}}}$ (its image being the space of functions $\overline{\mathcal{M}} \rightarrow \{W, NW\}$). Specifically, $s_n^2(s^1)$ is a function $s_{n|\mathcal{M}}^1 \rightarrow \{W, NW\}$ prescribing, for each $m \in \mathcal{M}_n$, whether or not to withdraw it. As there is no such decision, however, to be made on menus that have been offered with commitment, we impose the restriction

$$s_n^2(m, s_{-n}^1) = NW \quad \forall m \in s_{n|C}^1, \forall n \in \mathcal{N}, \forall s^1 \in S^1$$

To complete defining the firms' strategy space, let S_n^2 be the space of functionals $S^1 \rightarrow \{W, NW\}^{\overline{\mathcal{M}}}$, $S^2 = \prod_{n \in \mathcal{N}} S_n^2$, and $S_F = S^1 \times S^2$.

¹³Recall that our firms do observe at stage 3 the applications their contractual offers receive by the customers at stage 2. The very act of applying, therefore, for a given contract from a given menu institutes a signal a customer (the informed party) sends to the firms (the uninformed party). To fix ideas more precisely, imagine that there are two receivers in the well-known "beer-quiche" example from Cho and Kreps [9]. Suppose moreover that each receiver is in fact a drinks' vendor: one being the sole supplier of beer, while the other the sole supplier of quiche. To see that this corresponds to the typical subgame of stages 2-3 in our mechanism, in Figures 3-6 below, let "beer" and "quiche" stand, respectively, for the menu offered by the "entrant" and the "incumbent" (with "fight" and "acquiesce" standing, respectively, for "honor" and "withdraw").

Regarding the strategy of the typical customer at the beginning of stage 2, for risk-type h , this is given by a function $\alpha_h : S^1 \rightarrow \mathcal{A}$ which prescribes the contract to be chosen out of the available set

$$\mathcal{M}(s^1) = \cup_{n \in \mathcal{N}} s_n^1 | \mathcal{M}$$

in the subgame $\Gamma(s^1)$ for each history $s^1 \in S^1$. By the very definition of our mechanism, the contract in question is unique. Whenever it can be found in more than one menus, we will assume that each of these menus receives the same share of the total applications made for this contract by the risk-type h . For either risk-type, the strategy space is $S_A = \mathcal{A}^{S^1}$.

Equilibrium

To make predictions, we will use the pure-strategies version of the PBE in Maskin and Tirole [25]. This is basically a pure-strategies sequential equilibrium, but for the following caveat. The firms' contractual action space in stage 1 being infinite, the requirement for consistency of beliefs between stages 1-2 is not well-defined. For this reason we seek instead a vector of (pure) strategies (a profile of strategies for the firms and a strategy for each type of customer), and a vector of beliefs at each information set of the typical subgame $\Gamma(s^1)$ such that (i) the strategies are optimal under these beliefs (sequential rationality), and (ii) the belief at the beginning of $\Gamma(s^1)$ is the prior $\{\lambda, 1 - \lambda\}$ while the subsequent beliefs are everywhere (fully) consistent with the strategies.¹⁴ Under this notion of equilibrium and using the expression "honoring a menu" to mean that it is not withdrawn at stage 3, an equilibrium menu is such that (a) on the equilibrium path, it is honored at stage 3 and chosen by at least one risk-class of customers, and (b) there is no other menu that, if offered alongside the one in question, would expect strictly positive profits under some (fully) consistent profile of beliefs.

Formally, consider a given strategic profile $(s, \alpha_L, \alpha_H) \in S_F \times S_A \times S_A$. We first define the payoffs for the stage-2 subgame $\Gamma(s^1)$. Depicting the contract in the typical menu that is meant for the typical risk-type by the function $\mathbf{a}_h : \mathcal{M} \rightarrow \mathcal{A}$ with $\mathbf{a}_h = \mathbf{a}_h(m)$ for $m = \{\mathbf{a}_L, \mathbf{a}_H\}$, optimality requires that her best response is given by the function $\alpha_h : 2^{\mathcal{M}} \rightarrow \mathcal{A}$ where¹⁵

$$\alpha_h(M) \in \begin{cases} \arg \max_{\mathbf{a} \in \cup_{m \in M} \{\mathbf{a}_h(m)\}} U_h(\mathbf{a}) & \text{if } \exists m \in M : U_h(\mathbf{a}_h(m)) \geq \bar{u}_h \\ \{\mathbf{0}\} & \text{otherwise} \end{cases} \quad (1)$$

¹⁴The latter part of the statement in (ii) entails the existence of strategic profiles that are arbitrarily close to the equilibrium one and assign positive probability to every move on the game tree, such that the belief vectors that satisfy Bayes' rule for these strictly-positive strategic profiles are arbitrarily close to the equilibrium beliefs. In Maskin and Tirole [25] (see in particular their footnote 13), the requirement for consistency of beliefs between stages 1-2 is replaced by that of updating the prior to an interim belief via Bayes rule. In the present mechanism, the uninformed party being the one to make the proposals, no updating of beliefs takes place at the interim stage.

¹⁵Recall that we restrict attention to incentive-compatible menus while we let each risk-type choose at most one contract from the available set. There is no loss of generality thus in assuming that, out of a given menu m , the risk-type h chooses the contract $\mathbf{a}_h(m)$ - the one that is intended for her type. Notice also that no customer will ever choose any contract from a menu if she believes that the latter will be withdrawn. Under our normalization of the function u (recall footnote 10) any such menu is strictly dominated by the endowment lottery.

We will let also

$$\mathcal{M}_h(M) := \{m \in M : \mathbf{a}_h(m) = \alpha_h(M)\}, \quad M \in 2^{\mathcal{M}}$$

Turning next to the firms' payoff, observe that $\Gamma(s^1)$ is a signalling subgame in which the customers send signals to the firms by choosing a menu. For each $n \in \mathcal{N}$ and $s_n^1 \in S_n^1$, let $\lambda_n : s_{n|\mathcal{M}}^1 \times \{L, H\} \rightarrow [0, 1]$ be a belief function; $\{\lambda_{L_n}(m), \lambda_{H_n}(m)\}$ is the risk-type probability distribution the n th firm assigns to the pool of applicants for the menu m at stage 2. We need to distinguish of course between two kinds of information sets within $\Gamma(s^1)$. The first corresponds to offers that have been made with commitment, the respective information sets being singletons (see Figures 2-3 and the left panel of Figure 4). The second refers to offers that have been made without commitment, the respective information sets being non-singletons (see Figures 2, 4, and the right panel of Figure 3). At the typical information set, the firm's beliefs are given by

$$\lambda_{hn}(m) = \frac{\lambda_h \sigma_{hn}(m)}{\sum_{h \in \{L, H\}} \lambda_h \sigma_{hn}(m)} \quad h \in \{L, H\} \quad (2)$$

for some function $\sigma_{hn} : s_{n|\mathcal{M}}^1 \rightarrow \mathbb{R}_+$ and where $\lambda_L = \lambda = 1 - \lambda_H$; with the following restriction at the singleton information sets

$$\sigma_{hn}(m) = 1 \quad \forall h \in \{L, H\} \quad \forall m \in s_{n|C}^1, \quad \forall s_n^1 \in S_n^1, \quad \forall n \in \mathcal{N} \quad (3)$$

Suppose now that the n th firm is offering the collection $M_n \in 2^{\overline{\mathcal{M}}}$ while the typical from the other firms is honoring some $M_{n'} \in 2^{\overline{\mathcal{M}}}$. Letting

$$\mathcal{N}_h(\cup_{n \in \mathcal{N}} M_n) = \{n \in \mathcal{N} : M_n \cap \mathcal{M}_h(\cup_{n \in \mathcal{N}} M_n) \neq \emptyset\}$$

the payoff of the n th firm is given then by¹⁶

$$\begin{aligned} \Pi_n^W(M_n, \cup_{n \in \mathcal{N}} M_n, \alpha) &= 0 \\ \Pi_n^{NW}(M_n, \cup_{n \in \mathcal{N}} M_n, \alpha) &= \sum_{h \in \{L, H\}} \sum_{m \in M_n} \frac{\lambda_{hn}(m) \Pi_h(\mathbf{a}_h(m)) \mathbf{1}_{m \in \mathcal{M}_h(\cup_{n \in \mathcal{N}} M_n)}}{|M_n \cap \mathcal{M}_h(\cup_{n \in \mathcal{N}} M_n)| |\mathcal{N}_h(\cup_{n \in \mathcal{N}} M_n)|} \end{aligned}$$

if it withdraws and honors, respectively, each menu in M_n . Let now

$$\begin{aligned} \mathcal{M}(s_n) &= \{m \in s_{n|\mathcal{M}}^1 : s_n^2(m) = NW\} \\ \mathcal{M}(s) &= \cup_{n \in \mathcal{N}} \mathcal{M}(s_n) \quad \mathcal{M}(s_{-n}) = \cup_{n' \in \mathcal{N} \setminus \{n\}} \mathcal{M}(s_{n'}) \\ \Pi_n^{NC}(s_{n|\mathcal{M}}^1 \setminus M_n, s_{-n}, \alpha) &= \max_{M'_n \in 2^{s_{n|\mathcal{M}}^1 \setminus M_n}} \Pi_n^{NW}(M'_n, \mathcal{M}(s_{-n}) \cup M_n, \alpha) \end{aligned}$$

¹⁶The quantity $|\mathcal{N}_h(\cup_{n \in \mathcal{N}} M_n)|$ in the denominator of the typical term in the summation giving $\Pi_n^{NW}(M_n, \cup_{n \in \mathcal{N}} M_n, \alpha)$ ensures that the total profits/losses from a given risk-type are shared equally across those firms that honor her optimal contract. The quantity $|M_n \cap \mathcal{M}_h(\cup_{n \in \mathcal{N}} M_n)|$ ensures that, for each of these firms, the gain from offering this contract via multiple menus is not an accounting but a strategic one (in the sense of fending off or creating profitable deviations).

With respect to the decision of which menus to honor at the subgame $\Gamma (s^1)$ (out of the collections of menus on which it has not already committed at stage 1), the firm's best response function $s_n^{2*} : S_n^1 \times S_{-n} \times S_A \rightarrow 2^{s_n^1|_{\mathcal{M}}}$ is given by

$$s_n^{2*} \left(s_n^1|_{\mathcal{M}} \setminus M, s_{-n}, \alpha \right) \in \begin{cases} \arg \Pi_n \left(s_n^1|_{\mathcal{M}} \setminus M, s_{-n}, \alpha \right) & \text{if } \Pi_n \left(s_n^1|_{\mathcal{M}} \setminus M, s_{-n}, \alpha \right) \geq 0 \\ \{\emptyset\} & \text{otherwise} \end{cases}$$

for $M \in 2^{s_n^1|_{\mathcal{M}}}$ with $s_n^{2*}(m) = NW$ for each $m \in s_n^{2*} \left(s_n^1|_{\mathcal{M}} \setminus M, s_{-n}, \alpha \right)$.

Regarding the choice of menus on which to commit at stage 1, let

$$\begin{aligned} \Pi_n^C \left(M|s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right) &= \sum_{h \in \{L, H\}} |M \cap \mathcal{M}_h \left(\mathcal{M}(s_{-n}) \cup M \cup s_n^{2*} \left(s_n^1|_{\mathcal{M}} \setminus M, s_{-n}, \alpha \right) \right)|^{-1} \\ &\times \sum_{m \in M} \frac{\lambda_h \Pi_h(\mathbf{a}_h(m)) \mathbf{1}_{m \in \mathcal{M}_h \left(\mathcal{M}(s_{-n}) \cup M \cup s_n^{2*} \left(s_n^1|_{\mathcal{M}} \setminus M, s_{-n}, \alpha \right) \right)}}{|\mathcal{N}_h \left(\mathcal{M}(s_{-n}) \cup M \cup s_n^{2*} \left(s_n^1|_{\mathcal{M}} \setminus M, s_{-n}, \alpha \right) \right)|} \end{aligned}$$

again for $M \in 2^{s_n^1|_{\mathcal{M}}}$ (and where $\mathbf{1}_X$ is the indicator function of the set X). The firm's best response is given then by the functions $s_{n|C}^{1*} : s_n^1|_{\mathcal{M}} \times S_{-n} \times S_A \rightarrow 2^{s_n^1|_{\mathcal{M}}}$ with

$$s_{n|C}^{1*} \left(s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right) \in \begin{cases} \arg \max_{M \in 2^{s_n^1|_{\mathcal{M}}}} \Pi_n^C \left(M|s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right) & \text{if } \Pi_n^C \left(M|s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right) \geq 0 \\ \{\emptyset\} & \text{otherwise} \end{cases}$$

Letting finally

$$\begin{aligned} \Pi_n \left(s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right) &= \max \left\{ \begin{array}{l} \Pi_n^C \left(s_{n|C}^{1*} \left(s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right) |s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right), \\ \Pi_n^{NC} \left(s_n^1|_{\mathcal{M}} \setminus s_{n|C}^{1*} \left(s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right), s_{-n}, \alpha \right) \end{array} \right\} \\ \Pi_n \left(s^1, s_{-n}, \alpha \right) &= \max_{s_n^1|_{\mathcal{M}} \in \overline{\mathcal{M}}} \Pi_n \left(s_n^1|_{\mathcal{M}}, s_{-n}, \alpha \right) \end{aligned}$$

we are in position to define the equilibrium of Γ .

This is a profile of strategies (s^*, α^*) and belief functions $\{\lambda_{L_n}, \lambda_{H_n}\}_{n \in \mathcal{N}}$ such that $\alpha^*(\mathcal{M}(s^*)) = \alpha(\mathcal{M}(s^*))$ while for each $n \in \mathcal{N}$ we have $s_{n|C}^* \in \arg \Pi_n(s_{-n}^*, \alpha^*)$ if $\Pi_n(s_{-n}^*, \alpha^*) \geq 0$, with $c_n^*(m) = C$ [resp. $c_n^*(m) = NC$] if $m \in s_{n|C}^* \left(s_n^*|_{\mathcal{M}}, s_{-n}^*, \alpha^* \right)$ [resp. $m \notin s_{n|C}^* \left(s_n^*|_{\mathcal{M}}, s_{-n}^*, \alpha^* \right)$] and $s_n^{2*}(m) = NW$ [resp. $s_n^{2*}(m) = W$] if $m \in s_n^{2*} \left(s_n^*|_{\mathcal{M}} \setminus s_{n|C}^* \left(s_n^*|_{\mathcal{M}}, s_{-n}^*, \alpha^* \right) \right)$ [resp. $m \notin s_n^{2*} \left(s_n^*|_{\mathcal{M}} \setminus s_{n|C}^* \left(s_n^*|_{\mathcal{M}}, s_{-n}^*, \alpha^* \right) \right)$] while

$$\forall h \in \{L, H\}, \forall m \in s_n^{1*}|_{\mathcal{M}}, \quad \sigma_{hn}(m) = 1 \quad \text{if } m \in \mathcal{M}_h(\mathcal{M}(s^*)) \quad (4)$$

If $\Pi_n(s_{-n}^*, \alpha^*) < 0$, we can take any $s_{n|C}^* \in \overline{\mathcal{M}}$ with $c_n^*(m) = NC$ and $s_n^{2*}(m) = NW$ for all $m \in s_{n|C}^*$; or let $s_{n|C}^*$ consist of only the trivial menu $m_0 = \{\mathbf{0}, \mathbf{0}\}$ with $c_n^*(m_0) = C$ and $s_n^{2*}(m_0) = NW$.

Observe that the payoff of the typical firm on the equilibrium path is given by

$$\Pi_n(s^*, \alpha^*) = \frac{\lambda \Pi_L(\mathbf{a}_L(m_L))}{|\mathcal{N}_L(\mathcal{M}(s^*))|} \mathbf{1}_{m_L \in \mathcal{M}_L(\mathcal{M}(s^*))} + \frac{(1-\lambda) \Pi_H(\mathbf{a}_H(m_H))}{|\mathcal{N}_H(\mathcal{M}(s^*))|} \mathbf{1}_{m_H \in \mathcal{M}_H(\mathcal{M}(s^*))} \quad (5)$$

for some $(m_L, m_H) \in \mathcal{M}_L(\mathcal{M}(s_n^*)) \times \mathcal{M}_H(\mathcal{M}(s_n^*))$.

Efficient Insurance Provision

To obtain the MWS as the unique equilibrium allocation of our mechanism, we proceed in the following steps. We establish first that, in any candidate equilibrium of our mechanism, out of the set of menus offered at stage 1 on the equilibrium path, the most preferred one for the low-risk type must lie in the set $\{m \in \mathcal{M} : \mathbf{a}_L(m) = \bar{\mathbf{a}}_L, \mathbf{a}_H(m) \succsim \mathbf{a}_H^{**}\}$ (see Lemma A.1 in Appendix A and Claim 3.1 below). Of course, given free entry (see Lemma A.2 in Appendix A) and the fact that the MWS allocation solves uniquely the MWS problem, no such menu can end up being the equilibrium outcome unless it coincides with the MWS menu.

Next, we show that the MWS menu can be supported as equilibrium outcome (see Proposition 3.1 below). This leaves but the possibility for equilibrium scenarios in which menus from the set above are offered at stage 1 only to be withdrawn at stage 3 on the equilibrium path. That is, for equilibrium scenarios in which such menus are deployed as “latent” contractual offers: to be honoured in some off-equilibrium contingency, but never on the equilibrium path.¹⁷ Yet such scenarios do not survive the following equilibrium-selection criterion.

Definition Let $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ be an equilibrium profile for Γ . It is said to survive *commitment-irrelevance on the equilibrium path* if

$$\alpha_L(\mathcal{M}(s^*)) \sim_L \alpha_L(\mathcal{M}(s^{1*})) \tag{6}$$

Under this criterion, the MWS allocation is the unique equilibrium outcome of our mechanism (see Corollary 3.1 below).

With respect to the first step above, any equilibrium outcome $m^* = \{\alpha_L(\mathcal{M}(s^*)), \alpha_H(\mathcal{M}(s^*))\}$ must lie in \mathcal{M}_λ (recall again Lemma A.2). Given this, if $\mathbf{a}_L(m^*) \neq \bar{\mathbf{a}}_L$, we can find $\hat{m} \in \mathcal{M}_\lambda^+$ which constitutes a welfare improvement for the low-risk [resp. high-risk] type whenever $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(m^*)$ [resp. $\mathbf{a}_L(m^*) \succsim_L \bar{\mathbf{a}}_L$], and makes strictly-positive profits on the risk-type it attracts away from m^* (see Lemma A.3).

To see intuitively why such a menu exists, observe that we cannot have $\mathbf{a}_L(m^*) \succsim_L \bar{\mathbf{a}}_L$ unless $\bar{\mu} < 1$. Recall also that, whenever $\bar{\mu} = 1$ [resp. $\bar{\mu} < 1$], taking $\mu \in (\lambda, 1]$ [resp. $\mu \in (\bar{\mu}, 1]$] means that the low-risk [resp. high-risk] type is the profitable type at the $\text{IE}(\mu)$ -optimum while her welfare is strictly-increasing [resp. strictly-decreasing] in μ (see Lemmas 1.5 and 2.4 in the Online Appendix). Consequently, whenever $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(m^*)$ [resp. $\mathbf{a}_L(m^*) \succsim_L \bar{\mathbf{a}}_L$], letting μ approach $\bar{\mu}$ from below [resp. from above] guarantees in either case that the corresponding $\text{IE}(\mu)$ -optimum constitutes a welfare increase from m^* for the profitable type. To construct then \hat{m} it suffices to replace the contract \mathbf{a}_L^μ [resp. \mathbf{a}_H^μ] by some $\hat{\mathbf{a}}_L \in \mathcal{A}$ [resp. $\hat{\mathbf{a}}_H \in \mathcal{A}$] such that $\{\hat{\mathbf{a}}_L, \mathbf{a}_H^\mu\}$ [resp. $\{\mathbf{a}_L^\mu, \hat{\mathbf{a}}_H\}$] lies in the space \mathcal{M}_λ^+ (see Figure 1 for examples of this construction when the hypothetical equilibrium m^* is IE).

¹⁷Indeed, using the MWS menu itself as latent threat, the strategic scenario outlined in Proposition 3.1 supports as equilibrium outcome any allocation in \mathcal{M}_λ that is weakly dominated by the MWS one for the low-risk type (see Corollary A.1 in Appendix A).




Figure 1: Deviations $\hat{m} = \{\hat{\mathbf{a}}_L, \hat{\mathbf{a}}_H\} \in \mathcal{M}_\lambda^+$ against IE(μ)-optimal menus when $\mu \neq \bar{\mu}$: $w(\hat{\mathbf{a}}_L)$ lies in the shaded area while $\mathbf{w}(\hat{\mathbf{a}}_H) = \mathbf{w}_H^\mu = \mathbf{w}(\mathbf{a}_H^\mu)$; this gives $\Pi_H(\hat{\mathbf{a}}_H) < 0$ when $\mu < \bar{\mu}$ [left panel] and $\Pi_L(\hat{\mathbf{a}}_L) < 0$ when $\mu > \bar{\mu}$ [right panel]. Note: $FO_h^E = \{\mathbf{w}(\mathbf{a}) \in \mathcal{W} : \mathbf{a} \in \mathcal{A}, \Pi_h(\mathbf{a}) = 0\}$ is the iso-profit line through the endowment point $E = (W, W - d)$ that breaks even against risk-type h . Letting $\mathbf{a}^\mu \in \mathcal{A}$ given by $(1 - p_h)(a_{0h}^\mu - a_0^\mu) = p_h(a_{1h}^\mu - a_1^\mu)$ for $h \in \{L, H\}$ (see Lemma 2.7 in the Online Appendix), $FO_h^\mu = \{\mathbf{w}(\mathbf{a}) \in \mathcal{W} : \mathbf{a} \in \mathcal{A}, \Pi_h(\mathbf{a}) = \Pi_h(\mathbf{a}^\mu)\}$ is the iso-profit line through the point $\mathbf{w}(\mathbf{a}^\mu)$ that breaks even against risk-type h .

Notice now that, if no menu has been introduced at stage 1 that includes $\bar{\mathbf{a}}_L$, there are no menus on offer that could be used as threats (even latent ones) against \hat{m} . It follows, therefore, that a profitable deviation against m^* consists of offering \hat{m} with commitment. Being guaranteed then a strictly-better outcome, the profitable risk-type will choose \hat{m} at stage 2. Equally importantly, the latter remains strictly profitable even if the non-profitable type chooses it as well; hence, irrespective of the other firms' responses at stage 3.

Claim 3.1 $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ is an equilibrium profile for Γ only if $\alpha_L(\mathcal{M}(s^{1*})) = \bar{\mathbf{a}}_L$.

Proof. See Appendix A. ■

We show next that the MWS allocation $\bar{m} = \{\bar{\mathbf{a}}_L, \bar{\mathbf{a}}_H\}$ can be supported as equilibrium. To this end, observe first that \bar{m} must be offered without commitment when it differs from the RSW allocation. For withdrawal is the only way to fend off cream-skimming deviations ($\hat{m} \in \mathcal{M} : \mathbf{a}_L(\hat{m}) \succ_L \bar{\mathbf{a}}_L, \bar{\mathbf{a}}_H \succ_H \mathbf{a}_H(\hat{m})$) with commitment (see Figure 3). Indeed, since the MWS and RSW allocations differ only if $\bar{\mu} = 1$, the deviation cannot attract the low-risk type unless $\Pi(\hat{m}|\lambda) < 0$. And as the deviation is offered with commitment, the withdrawal of \bar{m} means that either risk-type will certainly choose the deviant menu, rendering the latter loss-making as it ends up attracting the entire population of customers (see the left panel of Figure 4). When the MWS and RSW allocations coincide on the other hand, which embeds the case $\bar{\mu} < 1$, whether or not the MWS menu is offered with commitment does not matter: any offering firm is indifferent between honoring or withdrawing




Figure 2: The opening stage (stage 1) of the game when attention is restricted to two firms, a deviant (the entrant) and one which offers the equilibrium menu (the incumbent). The nodes N.1-N.4 are the starting nodes for the signalling subgames of stages 2-3 (see Figures 3-4)

\bar{m} . Our strategic scenario has \bar{m} honoured in this case against all deviations.¹⁸ Without loss of generality, therefore, we may take the MWS menu to be offered without commitment in either case.

To account for cream-skimming deviations without commitment, we will illustrate the intuition behind our sequential equilibrium scenario restricting attention to deviations entailing single menus.¹⁹ For this, we must introduce some more notation. We will consider sequences of trembles $\{\epsilon_{kL}, \epsilon_{kH}\}_{k \in \mathbb{N}} \in (0, 1) \times (0, 1)$ with $\lim_{k \rightarrow \infty} (\epsilon_{kL}, \epsilon_{kH}) = (0, 0)$. These describe the firms' beliefs that risk-type h applies for \bar{m} and \hat{m} with probability $\sigma_{kh}(\bar{m}) = 1 - \epsilon_{kh}$ and $\sigma_{kh}(\hat{m}) = \epsilon_{kh}$. Given (2), the conditional beliefs are given by $\lambda_{hn}(m) = \lim_{k \rightarrow \infty} \lambda_{hn}^k(m)$ where $(n, m) \in \{(n^*, \bar{m}), (\hat{n}, \hat{m})\}$ and

$$\lambda_{Ln}^k(m) = \frac{\lambda_n^k \lambda}{\lambda_n^k \lambda + 1 - \lambda} = 1 - \lambda_{Hn}^k(m), \quad \lambda_n^k = \begin{cases} \frac{\epsilon_{kL}}{\epsilon_{kH}} & \text{if } n = \hat{n} \\ \frac{1 - \epsilon_{kL}}{1 - \epsilon_{kH}} & \text{if } n = n^* \end{cases} \quad h \in \{L, H\} \quad (7)$$

Moreover, for any menu $m \in \mathcal{M}$ and any belief $\{p, 1 - p\}$ in the simplex $\Delta([0, 1])$, we will view the expected profits

$$\Pi(m|p) = p\Pi_L(\mathbf{a}_L(m)) + (1 - p)\Pi_H(\mathbf{a}_H(m))$$

as a function of p , and define the solution

$$p(m) \in [0, 1] : \Pi(m|p(m)) = 0$$

The relation

$$\Pi(m|p) - \Pi(m|p(m)) = [p - p(m)] [\Pi_L(\mathbf{a}_L(m)) - \Pi_H(\mathbf{a}_H(m))]$$

¹⁸Recall Remark 1 in the preceding section and Figure 1 above. Letting $h, h' \in \{L, H\}$ with $h \neq h'$, cream-skimming deviations refer to menus $\hat{m} \in \mathcal{M}$: $\mathbf{a}_h(\hat{m}) \succsim_h \bar{\mathbf{a}}_h, \bar{\mathbf{a}}_{h'} \succ_{h'} \mathbf{a}_{h'}(\hat{m})$ where $h = L$ [resp. $h = H$] when $\bar{\mu} = 1$ [resp. $\bar{\mu} < 1$].

¹⁹As we have but two risk-types, the complete argument (see the proof of Proposition 3.1 in Appendix A) accounts for deviations entailing honoring up to two menus: m_L and m_H such that $\hat{m} = \{\mathbf{a}_L(m_L), \mathbf{a}_H(m_H)\}$ satisfies (IC_h) for either h . For this case replace the sequences $\{\epsilon_{kh}\}, \{\varepsilon_{kh}\}$ above by $\{\epsilon'_{kh}\}, \{\varepsilon'_{kh}\}$ for $\hat{m} = m_L$ and by $\{\epsilon''_{kh}\}, \{\varepsilon''_{kh}\}$ for $\hat{m} = m_H$. Replace also $\lambda_{n^*}^k = (1 - \epsilon_{kL}) / (1 - \epsilon_{kH})$ in (7) with $\lambda_{n^*}^k = (1 - \epsilon'_{kL} - \epsilon''_{kL}) / (1 - \epsilon'_{kH} - \epsilon''_{kH})$.




Figure 3: If offered with commitment, the MWS menu cannot fend off cream-skimming deviations that attract away the low-risk type. Let $\lambda = 1/2$, $U_L(\bar{\mathbf{a}}_L) = 1 = U_H(\bar{\mathbf{a}}_H)$, $2 = U_L(\hat{\mathbf{a}}_L) = 4U_H(\hat{\mathbf{a}}_H)$, and $\Pi_L(\bar{\mathbf{a}}_L) = 1 = -\Pi_H(\bar{\mathbf{a}}_H) = -\Pi_H(\hat{\mathbf{a}}_H) = 2\Pi_L(\hat{\mathbf{a}}_L)$. The unique PBE has the low-risk [resp. high-risk] type applying for the deviant [resp. MWS] menu ($\lambda_{LE}^k(\hat{m}) = 1 = 1 - \lambda_{LI}^k(\bar{m})$) and the deviant menu being honored, irrespective of whether it has been offered with [left panel] or without [right panel] commitment.

gives then the range of beliefs according to which the menu avoids losses:

$$\mathcal{P}(m) = \begin{cases} [0, 1] & \text{if } \Pi_L(\mathbf{a}_L(m)) = \Pi_H(\mathbf{a}_H(m)) \\ [p(m), 1] & \text{if } \Pi_L(\mathbf{a}_L(m)) > \Pi_H(\mathbf{a}_H(m)) \\ [0, p(m)] & \text{otherwise} \end{cases}$$

Remark 3. We have $\Pi(\bar{m}|\lambda) = 0$ while $\Pi_L(\bar{\mathbf{a}}_L) \geq \Pi_H(\bar{\mathbf{a}}_H)$ with equality if and only if the RSW allocation is IE (see Lemma 2.4 in the Online Appendix). Thus, if the MWS and RSW allocations coincide, we have $\mathcal{P}(\bar{m}) = [0, 1]$; otherwise, we have $\mathcal{P}(\bar{m}) = [\lambda, 1]$. Yet, $\lambda \in \mathcal{P}(\bar{m})$ in either case.

Observe now that, if $\bar{\mu} = 1$, then it must be $\Pi(\hat{m}|\lambda) < 0$ and thus $\lambda \notin \mathcal{P}(\hat{m})$. There are then two cases to consider: $\mathcal{P}(\hat{m}) = [p(\hat{m}), 1] \subset (\lambda, 1]$ and $\mathcal{P}(\hat{m}) = [0, p(\hat{m})] \subset [0, \lambda)$. In either case we can support the scenario in which \hat{m} gets withdrawn while \bar{m} is not, and either risk-type chooses the latter menu. If $\mathcal{P}(\hat{m}) \subset (\lambda, 1]$ (see the right panel of Figure 4 where $p(\hat{m}) = 2/3 > 1/2 = \lambda$), to ensure that \hat{m} is withdrawn let $\lim_{k \rightarrow \infty} \epsilon_{kL}/\epsilon_{kH} = 0$; in the limit, the deviant firm believes that its pool of applicants consists entirely of high-risk customers. If $\mathcal{P}(\hat{m}) \subset [0, \lambda)$, let $\epsilon_{kL}/\epsilon_{kH} = 1$; the deviant believes that the proportion of low-risk customers in its pool of applicants is exactly λ . In either case, moreover, we have $\epsilon_{kL} \leq \epsilon_{kH}$ (at least along a subsequence); that is, $\lambda_{n^*}^k \geq 1$ so that $\lambda_{Ln^*}^k(\bar{m}) > \lambda$ and $\lambda_{Hn^*}^k(\bar{m}) < 1 - \lambda$. As the trembles vanish, $\lambda_{Ln^*}^k(\bar{m})/\lambda_{Hn^*}^k(\bar{m})$ approaches $\lambda/(1 - \lambda)$ from above; in the limit, it becomes optimal to not withdraw \bar{m} - recall that $[\lambda, 1] \subseteq \mathcal{P}(\bar{m})$.

If $\bar{\mu} < 1$, on the other hand, then the MWS and RSW allocations coincide. In addition, since $\mathbf{a}_L(\hat{m}) \succ_L \mathbf{a}_L^* \succ_L \mathbf{a}_H^* \succ_H \mathbf{a}_H(\hat{m}) \succ_L \mathbf{a}_L(\hat{m})$, the menu $\{\mathbf{a}_L(\hat{m}), \mathbf{a}_H^*\}$ is separating. Clearly, the RSW allocation being IE, it cannot but be $\Pi_L(\mathbf{a}_L(\hat{m})) < 0$ (see the right panel of Figure 1). Let




Figure 4: The deviation consists of the same menu as in Figure 3. If the deviant menu offered with commitment [left panel], the unique PBE has the MWS menu withdrawn and either risk type choosing the deviant ($\lambda_{LE}(\hat{m}) = \lambda$). If the deviant menu is offered without commitment [right panel], the PBE has either type applying for the MWS menu ($\lambda_{LI}(\bar{m}) = \lambda$), which is honored while the deviant is withdrawn as loss-making ($\lambda_{LE}(\hat{m}) = 0$).

now $\lim_{k \rightarrow \infty} \epsilon_{kL}/\epsilon_{kH} = \infty$; in the limit, the deviant firm believes that the pool of applicants for \hat{m} consists entirely of low-risk customers (see the left panel of Figure 5). As for \bar{m} , the RSW allocation being safe, any offering firm is always indifferent between withdrawing it or not - whatever its belief about the risk-type distribution in the pool of applicants (recall that $\mathcal{P}(\bar{m}) = [0, 1]$ in this case).

Notice finally that deviations attracting both risk-types away from \bar{m} have $\mathbf{a}_L(\hat{m}) \succsim_L \bar{\mathbf{a}}_L$ while $\mathbf{a}_H(\hat{m}) \succsim_H \bar{\mathbf{a}}_H \succsim_H \mathbf{a}_H^{**}$; thus, $\Pi(\hat{m}|\lambda) < 0$ and the argument proceeds as above. Deviations attracting only the high-risk type have $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(\hat{m})$ and $\mathbf{a}_H(\hat{m}) \succ_H \bar{\mathbf{a}}_H \succ_H \mathbf{a}_H^{**}$; thus, $\Pi_H(\mathbf{a}_H(\hat{m})) < 0$ and not withdrawing \bar{m} renders them loss-making.

The equilibrium scenario just outlined can be described formally as follows.

Proposition 3.1 *The allocation $\{\bar{\mathbf{a}}_L, \bar{\mathbf{a}}_H\}$ is supported as the equilibrium outcome of Γ by the following (symmetric) strategy and belief profiles. Each firm $n \in \mathcal{N}$ chooses the action $s_{n|\mathcal{M}}^{1*} = \{\bar{m}\}$ with $c_n^*(\bar{m}) = NC$, and follows the strategy $s_n^{2*} : \{\bar{m}\} \rightarrow \{W, NW\}$ given by*

$$s_n^{2*}(\bar{m}) = \begin{cases} W & \text{if } \bar{m} \neq m^{**}, \exists (\hat{n}, \hat{m}) \in \mathcal{N} \setminus \{n\} \times s_{\hat{n}|C}^1 : \mathbf{a}_L(\hat{m}) \succsim_L \bar{\mathbf{a}}_L \\ NW & \text{otherwise} \end{cases}$$

Each customer follows the strategy in (1) while the firms' beliefs are as in (2)-(3).

Proof. See Appendix A. ■

It remains to show that the MWS is the unique equilibrium allocation under the selection-criterion in (6). Given though Claim 3.1, this is obvious.

supported by the deployment of “latent” menus.²⁰

A similar role, but with respect to the high-risk type and the MWS menu, is played also by the requirement that menus are the objects of trade. The MWS menu expects losses from the high-risk type whenever the MWS and RSW allocations differ (see Lemma 2.4 in the Online Appendix). If individual contracts were the objects of trade, on the equilibrium path, where all deviations are fended off and the MWS menu is left to separate the applicants by risk-type, it would be dominant strategy for the firm to withdraw only $\bar{\mathbf{a}}_H$ at stage 3.²¹ This would preclude though an equilibrium in the subgame of stages 2-3: the high-risk type’s optimal response at stage 2 would be to apply for $\bar{\mathbf{a}}_L$ instead of $\bar{\mathbf{a}}_H$, rendering the former a loss-making pooling contract.²²

In the context of the canonical insurance problem, the requirement that menus are the objects of trade is sine qua non for any mechanism with a stage in which the firms may withdraw their offers. For even when this stage precedes the one in which the customers choose, a slight modification of the preceding scenario destroys again the MWS equilibrium: on the equilibrium path, the dominant strategy would now be to withdraw $\bar{\mathbf{a}}_L$, forcing the low-risk type to choose $\bar{\mathbf{a}}_H$ and render it a strictly-profitable pooling contract (recall footnote 22).

A mechanism of this kind is analyzed in Netzer and Scheuer [30], where the time-line for the last two stages of the game in Hellwig [19] has been reversed: the customers choose at stage 3 while the firms may withdraw their offers at stage 2, albeit subject to paying a withdrawal cost. For sufficiently-small withdrawal costs, the mechanism supports the MWS allocation as the unique subgame-perfect Nash equilibrium (SPNE) outcome (see Proposition 2 in Netzer and Scheuer [30]). The MWS allocation is an equilibrium outcome also when the withdrawal costs are zero but in this case uniqueness fails: the set of equilibrium outcomes includes also any other allocation in \mathcal{M}_λ , when the MWS and RSW allocations coincide, and any other allocation in \mathcal{M}_λ that is strictly Pareto-dominated by the MWS allocation, otherwise (see Proposition 1 in Netzer and Scheuer [30]).

The authors embed the requirement that menus are the objects of trade within a stronger condition: a firm can only withdraw at stage 2 the entire set of contractual offers it made at

²⁰Recall that $\alpha_h^*(\mathcal{M}(s^{1*})) = \alpha_h(\mathcal{M}(s^*))$ for either h . And since $\mathcal{M}(s^*) \subseteq \mathcal{M}(s^{1*})$, the violation of (6) necessitates the existence of $m \in s_n^1 | \mathcal{M}$ for some $n \in \mathcal{N}$ such that $U_L(\mathbf{a}_L(m)) > \max_{m' \in \mathcal{M}(s^*)} U_L(\mathbf{a}_L(m'))$. But then the low-risk type would be better off in equilibrium if, other things being equal, s_n^1 prescribed that m is offered with commitment. Moreover, it must be the case that, on the equilibrium path, m gets introduced at stage 1 only to be withdrawn at stage 3. Clearly, its presence within the equilibrium scenario s^* is relevant only if the menu acts as “latent” threat - to be honored only against some deviation.

²¹To put it in the terminology of industrial organization (and to borrow from the relevant discussion in Asheim and Nilssen [3]), we restrict the extent to which firms can discriminate amongst customers. Our firms are allowed only withdrawals that do not discriminate customers on the basis of their contractual choices. They may engage in second-degree price discrimination (by offering non-linear pricing schemes) but not in third-degree price discrimination (to the extent that it would arise by segmenting the market according to the customers’ contractual choices). The latter restriction may originate in laws prohibiting third-degree price discrimination or in (so-called most-favoured-customer) clauses on the contracts themselves.

²²See Lemma 2.5(v) in the Online Appendix, and observe that $\Pi_H(\mathbf{a}) < \Pi_L(\mathbf{a})$ for any $\mathbf{a} \in \mathcal{A} \cap \mathbb{R}_+^2 \setminus \{\mathbf{0}\}$. It follows, therefore, that $\Pi(\mathbf{a}_L(m) | \lambda) < 0 < \Pi(\mathbf{a}_H(m) | \lambda)$ for any $m \in \mathcal{M}$: $\Pi(m | \lambda) = 0$. To apply this observation on equation (5), notice that, whenever the MWS and RSW menus differ, all active firms trade with both risk-types on the equilibrium path (see Lemma A.5).

stage 1, not individual offers from this set. And this highlights an important contrast, in economic as well as game-theoretic terms, with our mechanism. In economic terms, it means that a firm can withdraw an offer (a given insurance product in the current context) only by closing down completely. In game theoretic terms, it restricts the space of deviations under consideration.²³

Let the MWS and RSW allocations differ, and consider a deviation $s_{n|\mathcal{M}}^1 = \{\hat{m}, m\}$ against the MWS menu with \hat{m} a cream-skimming menu and $m \in \mathcal{M}_\lambda$. It is easy to see why Netzer and Scheuer do not let the deviant firm withdraw \hat{m} by itself. Facing \hat{m} , any firm that has entered \bar{m} at stage 1 cannot but withdraw it at stage 2. This would still leave though the deviant firm with strictly-positive profits: by withdrawing only \hat{m} , it would be left servicing the entire market with m . By contrast, such deviations do not affect at all our equilibrium scenario. Since $\bar{\mu} = 1$ when the MWS and RSW allocations differ, it cannot but be $\Pi(\hat{m}|\lambda) < 0$. As a result, if \hat{m} is introduced with commitment, withdrawing \bar{m} renders the deviation loss-making. And when \hat{m} gets introduced without commitment, our equilibrium scenario has \bar{m} being honored, not withdrawn (see the right panel of Figure 5 and Figure 6).²⁴




Figure 6: [Left panel] Deviations consisting of the menus $\hat{m} = \{\hat{\mathbf{a}}_L, \bar{\mathbf{a}}_H\}$ (with $\hat{\mathbf{a}}_L$ in the shaded area) and $m = \{\mathbf{a}, \mathbf{a}\}$ with $\mathbf{a} = (a_0^D, \bar{a}_{1L})$. [Right panel] Let $U_L(\mathbf{a}) = 0.8 = U_H(\mathbf{a})$ and $U_L(\hat{\mathbf{a}}_L) = 2$. The unique PBE has either risk type choosing the MWS menu ($\lambda_{LI}(\bar{m}) = \lambda$), which is honored while \hat{m} is withdrawn as loss-making ($\lambda_{LE}(\hat{m}) = 0$). Even though the deviant's strategy with respect to m is irrelevant, the payoffs depict it as honoring m .

This major difference between our mechanism and that in Netzer and Scheuer [30] stems from

²³In fact, strengthening in our mechanism the restriction that menus are the objects of trade to the requirement that only the entire set of contractual offers made at stage 1 may be withdrawn at stage 3 would ensure that the MWS is the unique equilibrium allocation, without the need for equilibrium selection: (6) would be trivially satisfied.

²⁴The MWS allocation being the IE(1)-optimum, $m \in \mathcal{M}_\lambda$ means that $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(m)$. If $\mathbf{a}_H(m) \succ_H \bar{\mathbf{a}}_H$, since $\bar{\mathbf{a}}_H \succ \mathbf{a}_H^*$, it cannot but be $\Pi_H(\mathbf{a}_H(m)) < 0$; offering then m with commitment renders it loss-making; offering m without commitment means that the deviant firm's best response is to subsequently withdraw it. If $\bar{\mathbf{a}}_H \succ_H \mathbf{a}_H(m)$, since then $\bar{\mathbf{a}}_h \succ_h \mathbf{a}_h(m)$ for either h , the deviant's strategy is inconsequential.

the latter’s reversal of the last two stages in Hellwig’s game. Having the customers as the last party to act renders the stages 2-3 subgame a screening game: when been called upon to act, either risk-type is certain of the delivery of her chosen contract; hence, either risk-type has always a (weakly) dominant strategy. This being of course common knowledge, the equilibrium outcome gets determined exclusively by the strategic interaction between the firms: a firm’s beliefs depict uncertainty about the actions of the other firms, not of the customers; as a result, the best response against a cream-skimming deviation cannot be but withdraw.²⁵

In our mechanism, this obtains only if the deviation has been offered with commitment. If it has not, the signaling subgame of stages 2-3 grants the strategic initiative to the insurance customers (the informed party). When called upon to act now, the optimal choice of either risk-type is not certain: it depends upon their belief about the likelihood that an offering firm withdraws a given menu. The customers’ uncertainty affects in turn the firm’s belief about the menu’s profitability. As a result, a cream-skimming menu gets withdrawn because the offering firm cannot be certain about the customers’ actions: it cannot be certain that its pool of applicants consists of only the low-risk type. The equilibrium outcome is determined now by strategic interaction not only between firms, but also between firms and customers.

Yet another difference between the two mechanisms has to do with the fact that the equilibrium scenario in Netzer and Scheuer [30] hinges crucially on the deployment of “latent contractual offers.” Depending on the equilibrium outcome to be supported, cream-skimming deviations are deterred because either the RSW menu or the MWS menu and the high-risk RSW contract are deployed as credible threats by “latent firms” (which in effect collude strategically with those offering the equilibrium menus).²⁶ In sharp contrast, our equilibrium-selection criterion in (6) precludes the use of latent menus.²⁷

Rather than modify the time-line of the game in Hellwig [19], the approach in Asheim and Nilssen [3] is to have the firms at stage 3 add rather than withdraw menus, subject to the condition

²⁵For a graphical representation of the game in Netzer and Scheuer [30] consider Figures 2-3 and the left panel of Figure 4. Replace the actions C and NC , respectively, by the actions NW and W , and the node N.4 by the payoff $(0, 0)$ in Figure 2. Delete also the NW action from the Incumbent [resp. Entrant] in the right panel of Figure 3 [resp. in the left panel of Figure 4].

²⁶The need for “latent firms,” not just “latent menus,” is noteworthy. For it shows that allowing firms to withdraw only their entire sets of contractual offers, not individual offers from their sets, restricts not just the space of possible deviations but the firms’ strategy space more generally. A firm offering the equilibrium menu in Netzer and Scheuer [30] is not allowed to use latent contractual offers on the equilibrium path (i.e., in the absence of deviations). In sharp contrast, it may well do so in our mechanism. This explains why the deployment of latent contractual offers in our mechanism allows also for equilibrium outcomes in \mathcal{M}_λ that are not Pareto-dominated by the MWS allocation, even when the MWS and RSW allocations differ (see Corollary A.1 in Appendix A). It also explains why withdrawal costs are of no use as means for equilibrium-selection in our mechanism.

²⁷Mimra and Wambach [27] extend the Netzer-Scheuer game along two dimensions: the firms are allowed to withdraw individual menus from their sets of contractual offers at stage 2, which entails a (potentially countably-infinite) sequence of withdrawal rounds until no firm wishes to make any more withdrawals. The authors’ aim is to provide a game-theoretic foundation for the notion of anticipatory foresight in Wilson [40]. The resulting set of equilibrium allocations is given by \mathcal{M}_λ . As in Netzer and Scheuer [30], the equilibrium scenario requires the deployment of latent contractual offers when the MWS and RSW allocations differ.

that their new offers (weakly) Pareto-dominate the original ones. This stage is then followed by another, the final stage, in which the customers choose from the revised set of menus. The MWS allocation emerges also here as the unique PBE; yet due to the fact that the firms' bargaining power is severely limited at stage 3: even though the firms may have already sorted their customers by type, cream-skimming offers are not allowed at this point.²⁸

Another strand of the recent literature considers the game-theoretic structure that supports the analysis in Rothschild and Stiglitz [35] (i.e., the restriction of the game in Hellwig [19] to only the first two stages), under appropriately chosen domains of contractual offers the firms may make.²⁹ In Inderst and Wambach [21], the firms face capacity constraints: each firm is limited in the number of contracts it can offer so that its customers may have to be rationed, while a rationed customer can search for another firm with free capacity but only at a cost. The model supports always the RSW allocation as equilibrium (even when this is not possible in Rothschild and Stiglitz [35] - see our Claim 4.3 below), as long as the search costs are chosen appropriately (to make any deviation face a pool of applicants whose risk-distribution is sufficiently less favorable than the prior to render it loss-making; see assumptions A.1-A.2 in Inderst and Wambach [21]).

The issue of existence of equilibrium in Rothschild and Stiglitz [35] is addressed also in Picard [31], with menus of contracts in which customers share the profits or losses of the insurer. When no equilibrium exists in Rothschild and Stiglitz [35] full participation by the customers in the profit or losses of the firms (ownership with unlimited liability) sustains the MWS menu as a SPNE (within a dense SPNE set where the MWS is the only IE-optimal allocation). The equilibrium argument rests upon the fact that the losses the MWS menu makes after a cream-skimming deviation will have

²⁸The authors' interpretation of the condition that only (weakly) Pareto-dominant new contractual offers are allowed at stage 3 is that of "non-discriminating renegotiation" between the firms and their customers. Even though much more restrictive in the kind of deviations it allows for, on the equilibrium path "non-discriminating renegotiation" has the same effect as our condition that menus are the object of trade: to ensure that the high-risk MWS contract will not be withdrawn. Even on the equilibrium path, however, there is a significant difference between the two conditions. It would take but expositional changes in our analysis if we were to allow the firms to withdraw individual contracts from menus, but give them also a two-dimensional commitment option: to pre-commit on a given menu at stage 1, if they so wish, upon either the delivery of one of its constituent contracts or the "non-discriminating withdrawal" of the menu itself (recall footnote 21). The former dimension suffices for the deviations that deliver Lemma A.3. The latter would be optimally chosen on the equilibrium path of the scenario in Proposition 3.1. In sharp contrast to "non-discriminating renegotiation," "non-discriminating withdrawals" can be viewed as an endogenously-emerging optimal choice by the firms themselves.

²⁹In a methodologically different approach, von Siemens and Kosfeld [38] impose instead conditions on the preferences of the informed player. They assume that the low-risk customers prefer to be pooled on a contract with their own rather than the high-risk type. This (one-sided) externality facilitates sorting and supports multiple separating allocations as equilibria (see also von Siemens and Kosfeld [37] for the application of this approach in the screening version of the standard job-market model with adverse selection). Yet another methodology is deployed by Ania et al. [2] where the dynamics of the market in Rothschild and Stiglitz [35] are modeled using evolutionary game theory. Specifically, the insurers are assumed to be boundedly rational in their strategic behavior: they limit themselves to imitating profitable contractual offers by other firms, and experimenting (in a particular way) with their own offers. The RSW allocation emerges as the unique log-run equilibrium outcome if experimentation is sufficiently limited. However, the result is not robust to even small improvements in firms' strategic sophistication (see Remark 3 and footnote 15 in Section 4 of Ania et al. [2]).

to be incurred by the high-risk customers, unless they also choose the deviating offer (rendering it loss-making).

4.1 Exit in Bertrand competition

By not restricting the set of contracts from or the way in which the firms make their offers, nor changing the timing under which these offers are made in Hellwig [19], the present study exhibits another fundamental difference compared to the literature: it embeds Hellwig's analysis as well as that in Rothschild and Stiglitz [35]. And by doing so, it sheds light in important aspects of Bertrand competition under asymmetric information that seem to have been ignored by the two papers - causing confusion in the subsequent literature.

To view the results in Hellwig [19] from the present perspective, set $\mathcal{C} = \{NC\}$; remove that is the pre-commitment option from the firms' decision at stage 1. To compare with the unrestricted version of Γ , recall that the equilibrium scenario in Proposition 3.1 calls for the MWS menu to be withdrawn only against cream-skimming deviant menus introduced with commitment. When the firms cannot commit to their offers, this part of the equilibrium scenario becomes irrelevant. And so does the potential use of latent menus to support equilibria; hence, also the requirement for commitment-irrelevance. The entire \mathcal{M}_λ becomes now the set of equilibrium allocations. It can be supported by an equilibrium scenario that has the respective menu always honored at stage 3.

Claim 4.1 *Set $\mathcal{C} = \{NC\}$ in Γ . The set of equilibrium allocations is \mathcal{M}_λ . Each $m \in \mathcal{M}_\lambda$ can be supported as equilibrium outcome by the following (symmetric) equilibrium strategy and belief profiles: each firm $n \in \mathcal{N}$ chooses the action $s_n^{1*} = (m, NC)$ and the strategy $s_n^{2*} : S^1 \rightarrow \{NW\}$, each customer follows the strategy in (1), while the associated beliefs for the firms are as in (2).*

Proof. See Appendix B. ■

The set of equilibrium allocations gets reduced if we restrict attention to allocations that give the high-risk type (the source of the adverse selection problem) at least her first-best welfare: i.e., equilibrium profiles $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ such that $\alpha_H(\mathcal{M}(s^*)) \succeq_H \mathbf{a}_H^{**}$. Needless to say, in practical terms, for this type to receive her first-best allocation it suffices to simply reveal her type; in game-theoretic terms, it suffices that some (zeroth) firm includes the (safe) menu $\{\mathbf{a}_H^{**}, \mathbf{a}_H^{**}\}$ in its collection of offers at stage 1 with the plan to not withdraw it at stage 3 irrespectively of the history of play at that point. Guaranteeing the high-risk type her first-best welfare precludes her from being profitable in equilibrium; necessitating in turn that the low-risk is the profitable type (see Lemma B.1). In this case, the set of equilibrium allocations in Claim 4.1 shrinks from \mathcal{M}_λ to the following subset

$$\mathcal{M}' = \{m \in \mathcal{M}_\lambda : \Pi_H(\mathbf{a}_H(m)) \leq 0 \leq \Pi_L(\mathbf{a}_L(m)), \mathbf{a}_H(m) \succeq_H \mathbf{a}_H^{**}\}$$

Of course, to obtain the exact setting in Hellwig [19] we must rule out cross-subsidization. We must restrict that is the space of contractual offers to singleton (i.e., pooling contracts) or *safe*

menus:

$$\mathcal{M}'' = \{m \in \mathcal{M} : \mathbf{a}_L(m) = \mathbf{a}_H(m) \vee \Pi_h(\mathbf{a}_h(m)) \geq 0 \forall h \in \{L, H\}\}$$

By Claim 4.1, the set of equilibrium allocations reduces now to $\mathcal{M}_\lambda \cap \mathcal{M}''$ or $\mathcal{M}' \cap \mathcal{M}''$ - depending of course on whether or not we wish to guarantee the high-risk type her first-best welfare. This illustrates the importance of cross-subsidization in Bertrand competition when firms cannot commit to not exit. The set $\mathcal{M}_\lambda \cap \mathcal{M}''$ includes at most two IE allocations, the pooling contract IE-(λ) and the RSW allocation when it is IE. By contrast, \mathcal{M}' includes the entire IE frontier when $\bar{\mu} = 1$, and the collection of IE(μ)-optima with $\mu \in [0, \bar{\mu}]$ otherwise (see Lemma 2.4 in the Online Appendix).

To view next the results in Rothschild and Stiglitz [35] from the perspective of our mechanism, set $\mathcal{C} = \{C\}$; remove that is the option of not committing from the firms' decision at stage 1. To compare again with the unrestricted version of Γ , observe that (6) holds now trivially (recall footnote 20). As a result, the MWS menu remains the unique candidate equilibrium allocation (recall Claim 3.1). It is easy to show moreover that the equilibrium allocation must break even on a contract-by-contract basis (see Lemma B.2). Clearly, when firms cannot withdraw their offers, the unique candidate equilibrium allocation is the MWS one when it coincides with the RSW allocation.

Observe also that no contract is able to attract only one risk-type away from the RSW allocation and avoid losses doing so (see Lemma 2.6 in the Online Appendix). Hence credible challenges may come only from menus $\hat{m} \in \mathcal{M}$ that attract both types away from the RSW allocation (i.e., $\mathbf{a}_h(\hat{m}) \succ_h \mathbf{a}_h^*$ for either h) and are strictly-profitable doing so (i.e., $\Pi(\hat{m}|\lambda) > 0$). Yet there are no such menus when the RSW allocation coincides with the MWS one. To summarize these observations more formally recall that the MWS and RSW allocations coincide when and only when the latter is IE (see Claim 3.1 in the Online Appendix).³⁰

Claim 4.2 *Set $\mathcal{C} = \{C\}$ in Γ . The RSW is the unique candidate equilibrium allocation. It is supported as equilibrium if and only if it is IE.*

To match exactly the setting in Rothschild and Stiglitz [35] we must restrict again the space of contractual offers to \mathcal{M}'' . The RSW menu being safe, it remains the unique candidate equilibrium allocation. However, as pointed out in Rothschild and Stiglitz [35] (albeit heuristically), the RSW allocation can be supported as equilibrium now if and only if λ is small enough for the market fair-odds line $\{\mathbf{a} \in \mathcal{A} : \Pi(\{\mathbf{a}, \mathbf{a}\}|\lambda) = 0\}$ to lie below the indifference curve of the low-risk type through \mathbf{a}_L^{**} .

As it turns out, the latter condition is not equivalent with the RSW allocation being IE. Letting \mathcal{U} denote the space of strictly-increasing, strictly-concave, and twice-continuously differentiable functions $[W - d - \bar{W}, W + \bar{W}] \rightarrow \mathbb{R}$, we can consider $\mathcal{U} \times (0, 1) \times (0, 1) \times (0, 1)$ as the space of the parameter-vectors (u, p_L, p_H, λ) for the canonical insurance problem. And \mathcal{U} being a complete

³⁰See also Dosis [11] for a direct proof of Claim 4.2 in a canonical insurance market with many risk-types albeit only two firms.

metric space under the sup-norm, we may define a measure on its σ -algebra of Borel sets along with the Borel measure on $(0, 1)$. We can show then that, over a non-null subset of the parameters' space, the RSW allocation is supported now as equilibrium even though it is not IE.

Claim 4.3 *Set $\mathcal{C} = \{C\}$ and $\mathcal{M} = \mathcal{M}''$ in Γ . The RSW is the unique candidate equilibrium allocation. It is supported as equilibrium if and only if*

$$U_L(\mathbf{a}_L^{**}) \geq \max_{\mathbf{a} \in \mathcal{A}: \Pi(\{\mathbf{a}\}|\lambda) \geq 0} U_L(\mathbf{a}) \quad (8)$$

There exist moreover $\lambda', \lambda'' \in (0, 1)$ with $\lambda' < \lambda''$ such that, everywhere in the space $\mathcal{U} \times (0, 1) \times (0, 1) \times (\lambda', \lambda'')$ of parameters for the canonical insurance problem, condition (8) is satisfied while the RSW allocation is not IE.

Proof. See Appendix B. ■

Cross-examining Claims 4.2-4.3 reveals the role of cross-subsidization when firms cannot exit the Bertrand-competitive market under study: it ensures that the equilibrium outcome will be efficient, at the expense though of making it harder for the market to achieve equilibrium. Intuitively speaking, and irrespective of whether or not we allow for cross-subsidization, if firms have to commit upon their offers they cannot but choose a safe menu on the equilibrium path (recall Lemma B.2). Equally importantly, withdrawal no longer an option, this menu can be supported as equilibrium only in the absence of profitable deviations - which under free-entry must attract away the low-risk type, the only profitable type (see Lemma B.1). Clearly, the only candidate equilibrium allocation is the RSW one, the most preferred for the low-risk type amongst the safe allocations.

Whether or not it can be supported as equilibrium depends on whether or not there exist profitable deviations. In the absence of cross-subsidization, such deviations can come only from singleton menus (pooling contracts): to survive as equilibrium, the RSW allocation must maximize the welfare of the low-risk type amongst such menus. By contrast, when firms are allowed to cross-subsidize between contracts, the set of profitable deviations is considerably larger: to survive now, the RSW allocation must maximize the welfare of the low-risk type within the space \mathcal{M}_λ .

In the light of this discussion, we can trace the limiting framework under which the notion of competitive equilibrium under adverse selection was examined in Rothschild and Stiglitz [35] but also in Hellwig [19]. By restricting attention to allocations in \mathcal{M}'' , the two seminal studies led the subsequent literature to endorse the premise that the efficiency of equilibrium is fundamentally related to its existence, and that either becomes an issue precisely because the contractual offers are restricted in \mathcal{M}'' .³¹ Yet this restriction is in either model of only minor consequence by itself: even though achieving efficiency is harder in the absence of cross-subsidization, it is not by any means guaranteed in its presence either.³²

³¹See for instance the relevant discussions in Mimra and Wambach [26] (pp. 140), Netzer and Scheuer [30] (pp. 1), von Siemens and Kosfeld [37] (pp. 182-3)-[38] (pp. 3-4), Ania et al. [2] (pp. 154), or Inderst and Wambach (pp. 1982).

³²Claim 3.1 in the Online Appendix attests to how strong is the condition that the RSW allocation is IE. Notice also that the IE frontier is but a subset of negligible relative size within \mathcal{M}' .

The real issue between the two models has to do instead with their respective approach towards market exit. In Rothschild and Stiglitz [35], by prohibiting firms to exit, the model guarantees efficiency of equilibrium at the expense of not doing so for existence. In Hellwig [19], by not allowing firms to commit to not exit, the model ensures existence of equilibrium at the expense of not doing so for efficiency. The present study suggests that the real desideratum for Bertrand competition under adverse selection is to simultaneously enable firms to exit but also publicly-commit to not do so.

4.2 A neutrally-optimal mechanism

Refocusing the analysis on the very design of this well-known competitive environment calls in turn for a comparison with two of the most seminal studies in mechanism design: Maskin and Tirole [25] and Myerson [29]. To fix ideas, it should be noted that throughout the present paper we have referred to the notions of (direct) mechanism and revelation principle having in mind but the standard definitions (see for instance Section 2 in Myerson [29]). With respect to either notion, the firms' option to withdraw at stage 3 can be interpreted explicitly as part of a direct mechanism: for any $m \in \mathcal{M}$ and for either risk-type, the action $((m, NC), W)$ maps to the trivial menu $\{\mathbf{0}, \mathbf{0}\}$. In this sense, the perceived lack of commitment by a firm to make good on its stage-1 offer refers to elements of strategy, not mechanism design. Choosing whether or not to publicly-commit to m amounts to choosing at stage 1 between committing to $\{m\}$ or committing to $\{m, \{\mathbf{0}, \mathbf{0}\}\}$ - where opting to commit to a doubleton of menus postpones for stage 3 the choice of which menu to deliver eventually as allocation.

Interpreting the extent of ex-ante commitment as an element of strategy brings us in line with the analysis in Maskin and Tirole [25], which also considers the game-form in Hellwig [19] yet under a significant generalization of the notion of contractual arrangement. A “contract” is an actual mechanism: it specifies a game to be played between two parties, the set of possible actions for each, and an allocation for each pair of their strategies. More precisely, the grand game begins with stage 1 where at least two uniformed parties (UPs) propose simultaneously “contracts” to an informed party (IP). The latter responds at stage 2: if she accepts a proposal, the corresponding game is played out and each party receives the respective outcome at stage 3; otherwise, each party gets its reservation payoff.³³

Even in the context of the canonical insurance problem, the ensuing set of equilibrium outcomes turns out to be very large. Any allocation in \mathcal{M}_λ that satisfies the profit constraint with equality can be supported as equilibrium.³⁴ The supporting strategic scenario prescribes that, following a

³³The main focus in Maskin and Tirole [25] (see Sections 4-6) is on a signaling adaptation of the game-form in Hellwig [19]. Nonetheless, to facilitate comparison with our analysis, we are referring here to their screening approach in Section 7. This entails no loss of generality: the main message of Section 7 is that the equilibrium outcomes under screening and signalling remain essentially the same, as long as the out-of-equilibrium actions and beliefs are left unrestricted.

³⁴See Proposition 12 as well as condition (iv) of Proposition 7 in Maskin and Tirole [25]. That the condition is satisfied in the canonical insurance problem follows immediately from Claim 2.1 in the Online Appendix.

strictly-profitable deviation by another UP, the outcome of the equilibrium mechanism would be an allocation that all IP-types prefer strictly to that of the deviant. Yet this allocation is but a latent part of the equilibrium scenario, which explains why the analysis in Maskin and Tirole [25] seems at odds with our result in Corollary 3.1: our equilibrium-selection criterion rules out the use of latent contractual offers on the equilibrium path. In conjunction with endogenous commitment, this allows our mechanism to restrict the market participants' out-of-equilibrium beliefs appropriately so as to deliver the MWS allocation as the unique equilibrium outcome.

An even sharper difference emerges when the Maskin-Tirole grand game gets modified into one in which the two parties alternate in making proposals. The IP being now able to exert influence on the mechanism design, the equilibrium set consists of the allocations in \mathcal{M}_λ that weakly Pareto-dominate the RSW one (see Proposition 13 in Maskin and Tirole [25]). And this remains the equilibrium set when the original grand game becomes instead a signalling one, with the IP making the proposals at stage 1 (see Proposition 6 in Maskin and Tirole [25]). Needless to say, the equilibrium set in question consists of only the MWS allocation if and only if the RSW allocation is IE.

In the canonical insurance problem, however, the MWS allocation is the only outcome the IP should aim for when she is called upon to propose mechanisms. This follows from the study in Myerson [29], which considers the signalling grand game above but for the generalization that all parties may be informed. This paper brought attention to the *core mechanisms* and characterized the *neutral optima* among them. The latter are core mechanisms that can be supported as sequential equilibria of the grand game while forming the smallest class of mechanisms that satisfy four fundamental axioms of mechanism selection.

Restricting attention to deterministic mechanisms (i.e., allocations), in the canonical insurance problem the notion of core is defined as follows (for the general definition, see Sections 3 and 6 in Myerson [29]). For an allocation to be core, there should not exist another allocation that would be (i) strictly preferred by at least one risk-type, and (ii) individually-rational for the firms given the information revealed to them when a type who strictly prefers the new allocation proposes it alone or when both types propose it. Formally, $m \in \mathcal{M}$ is a core allocation if there does not exist $\hat{m} \in \mathcal{M}_\lambda$ such that $\emptyset \neq T = \{h \in \{L, H\} : \mathbf{a}_h(\hat{m}) \succ_h \mathbf{a}_h(m)\}$ while $\Pi_h(\mathbf{a}_h(\hat{m})) \geq 0$ for all $h \in T$. By Lemma A.3, the MWS is the unique candidate core allocation. That it is in fact a core allocation, and thus the unique neutral optimum, follows immediately (see Remark V in the Online Appendix).

The theoretical import of a mechanism that delivers the unique neutral optimum as the unique equilibrium outcome cannot be overstated. The fact that our analysis sits comfortably within the realm of incentive-compatible direct mechanisms and the revelation principle begs the question as to how it fairs against generalized (direct as well as non-direct) mechanisms. In the canonical insurance problem, the type-spaces of the principal and the agent being a singleton and a doubleton, respectively, the requirement that the mechanism is incentive compatible reduces to our (IR_F) and (IC_h) constraints (see condition (2.6) in Myerson [29]). One could re-interpret therefore our menu space \mathcal{M} as the space of incentive-compatible mechanisms: each mechanism giving the

corresponding menu as its Bayesian-Nash equilibrium (BNE) allocation, and the customers' action being to choose between the two contracts in the given menu. Under this interpretation our result can be viewed as follows. By delivering the unique neutral optimum as the unique equilibrium outcome, our mechanism is a sequential equilibrium of the generalized-mechanism selection game for the canonical insurance problem (see Theorem 5 in Myerson [29]). Which is to say that every generalized mechanism for the canonical insurance problem has at least one BNE that is weakly dominated by ours for either risk-type (see Section 5 - especially inequality (5.7) - in Myerson [29]).

Equally noteworthy is the import of our analysis with respect to illustrating fundamental notions in Myerson [29]. To support the unique neutral optimum as the unique equilibrium outcome, the requisite restriction in the out-of-equilibrium beliefs obtains because we allow the firms to publicly pre-commit at stage 1, if they so wish, upon the delivery of their offers at stage 3 irrespectively of the history of play at that point. This can be viewed as an example of Myerson's *public actions*: enforceable decisions individual players can publicly-commit themselves to carry out, even if they may turn out ex-post to be harmful to themselves or others. Our notion of endogenous commitment demonstrates the crucial (yet abstract) role the set of public actions commands in Myerson [29] when it comes to establishing the very existence of neutral optima.

Of course, we are by no means the first to highlight the role of endogenous public commitment, nor that of the required supporting institutional structure. The idea of market environments that allow for commitment as endogenously-emerging strategic choice exists already in the literature. Nonetheless, a fundamental feature of the equilibrium scenario distinguishes our mechanism sharply: public commitment is called upon only off-equilibrium, never on the equilibrium path.³⁵ And this becomes the more important when it comes to applications of the canonical insurance model: viewing the equilibrium scenario as the predicted path of play, the fact that commitment is not chosen along this path goes a long way in alleviating the need of a motivational story for the implicitly-assumed commitment technology.

5 Concluding Remarks

The simplest and most obvious approach to studying markets with incomplete information is to consider the case of one-sided information: one side of the market has private information about some relevant characteristic of the good to be traded. In the context of the accident-insurance

³⁵To point out but a few examples, Caruana and Einav [10] consider a dynamic game in which players can change their previously-announced actions but at a cost. Commitment here is the default strategic option, and the authors show that it remains almost always the option of choice on the equilibrium path. Equally important is the role commitment plays on the equilibrium path in Netzer and Scheuer [30]. By contrast, Baker et al. [5] allow firms to choose between contracts that are binding (because they are based on objectively-verifiable even though poor performance measures), and contracts that are not (as they are written upon subjective but non-verifiable indicators). Quite similarly, Bernheim and Whinston [7] analyse binding versus non-binding contractual agreements as the choice between complete and incomplete contracts. As models of endogenous commitment, however, neither paper relates to the canonical insurance problem - where one assumes (among other things) contractual completeness in order to focus on the issue of adverse selection.

market in Rothschild and Stiglitz [35], the present study approaches a central issue that has been left open in the literature from a new perspective. Our analysis suggests that the lack of efficient outcomes in competitive markets under adverse selection may not be due to the presence of private, but rather due to the absence of public information. More precisely, due to the lack of institutions that facilitate public commitment by the insurance suppliers - to deliver on the offers customers have applied for via “pre-approved” forms, and to not discriminate by risk-type when rejecting applications via standard forms.

Our mechanism augments the game in Hellwig [19] along two dimensions, an accounting and a strategic. Our firms can subsidize their net income across contractual offers via the deployment of menus. They can also choose though to publicly pre-commit upon the future delivery of their menus. In the canonical insurance problem, this type of endogenous commitment plays the same strategic role as that of the public actions in Myerson [29]. As a result, our mechanism can deliver the unique neutral optimum as the unique equilibrium outcome.

In the context of the canonical insurance problem, our approach demonstrates the operational content of Myerson’s seminal argument for the existence of neutral optima, an ingenious though entirely abstract method of restricting the out-of-equilibrium beliefs appropriately. In this sense, the present paper suggests a rather unique in the current literature application of fundamental theoretical insight. In a more intuitive sense, the present mechanism can be seen as an optimality benchmark for the evaluation of other accident-insurance models in the literature that may offer closer approximations to real-world market settings.

References

- [1] Ales L. and P. Maziero (2012): “Adverse Selection and Non-exclusive Contracts,” *Tepper School of Business Working Paper No. 574* (available at <http://repository.cmu.edu/tepper/574/>).
- [2] Ania A.B., Tröger T., and A. Wambach (2002): “An Evolutionary Analysis of Insurance markets with Adverse Selection,” *Games and Economic Behavior*, 40:153-84.
- [3] Asheim G. and T. Nilssen (1996): “Non-discriminating Renegotiation in a Competitive Insurance Market,” *European Economic Review*, 40:1717-36.
- [4] Attar A., Marriotti T, and F. Salanié (2014): “Nonexclusive Competition under Adverse Selection,” *Theoretical Economics*, 9:1-40.
- [5] Baker G., Gibbons R., and K.J. Murphy (1994): “Subjective Performance Measures in Optimal Incentive Contracts,” *Quarterly Journal of Economics*, 109:1125-56.
- [6] Banks J.S. and J. Sobel (1987): “Equilibrium Selection in Signalling Games,” *Econometrica*, 55:647-61.

- [7] Bernheim D.B. and M.D. Whinston (1998): “Incomplete Contracts and Strategic Ambiguity,” *American Economic Review*, 88:902-32.
- [8] Bisin A. and P. Gottardi (2006): “Efficient Competitive Equilibria with Adverse Selection,” *Journal of Political Economy*, 114:485-516.
- [9] Cho I.K. and D. Kreps (1987): “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 102:179-221.
- [10] Caruana G. and L. Einav (2008): “A Theory of Endogenous Commitment,” *The Review of Economic Studies*, 75:99-116.
- [11] Dosis A. (2017): “Nash Equilibrium in Competitive Insurance,” *Economic Letters*, 152:5-8.
- [12] Dubey P. and J. Geanakoplos (2002): “Competitive Pooling: Rothschild-Stiglitz Reconsidered,” *Quarterly Journal of Economics*, 117:1529-70.
- [13] Dubey P., J. Geanakoplos, and M. Shubik (2005): “Default and Punishment in General Equilibrium,” *Econometrica*, 73:1-37.
- [14] Engers M. and L. Fernandez (1987): “Market Equilibrium with Hidden Knowledge and Self-selection,” *Econometrica*, 55:425-39.
- [15] Gale D. (1996): “Equilibria and Pareto Optima of Markets with Adverse Selection,” *Economic Theory*, 7:207-35.
- [16] Gale D. (1992): “A Walrasian Theory of Markets with Adverse Selection,” *The Review of Economic Studies*, 59:229-55.
- [17] Grossman H.I. (1979): “Adverse selection, Disassembling, and Competitive Equilibrium,” *The Bell Journal of Economics*, 10:336-43.
- [18] Guerrieri V., R. Shimer, and R. Wright (2010): “Adverse selection in Competitive Search Equilibrium,” *Econometrica*, 78:1823-62.
- [19] Hellwig M. (1987): “Some Recent Developments in the Theory of Competition in Markets with Adverse Selection,” *European Economic Review*, 31:319-25.
- [20] Holmstrom B. and R.B. Myerson (1983): “Efficient and Durable Decision Rules with Incomplete Information,” *Econometrica*, 51:1799-819.
- [21] Inderst R. and A. Wambach (2001): “Competitive Insurance Markets under Adverse Selection and Capacity Constraints,” *European Economic Review*, 45:1981-92.
- [22] Kohlberg E. and J.F. Mertens (1986): “On the Strategic Stability of Equilibria,” *Econometrica*, 54:1003-37.

- [23] Martin A. (2007): “On Rothschild-Stiglitz as Competitive Pooling,” *Economic Theory*, 3:371-86.
- [24] Mas-Colell A., Whinston M.D., and J.R. Green: *Microeconomic Theory*, Oxford University Press (1995).
- [25] Maskin E. and J. Tirole (1992): “The Principal-Agent Relationship with an Informed Principal, II: Common Values,” *Econometrica*, 60:1-42.
- [26] Mimra W. and A. Wambach (2014): “New Developments in the Theory of Adverse Selection in Competitive Insurance,” *The Geneva Risk and Insurance Review*, 39:136-52.
- [27] Mimra W. and A. Wambach (2018): “Contract Withdrawals and Equilibrium in Competitive Insurance Markets with Adverse Selection,” *Economic Theory* (forthcoming).
- [28] Miyazaki H. (1977): “The Rat Race and Internal Labor Markets,” *The Bell Journal of Economics*, 8:394-418.
- [29] Myerson R.B. (1983): “Mechanism Design by an Informed Principal,” *Econometrica*, 51:1767-97.
- [30] Netzer N. and F. Scheuer (2014): “A Game Theoretic Approach of Competitive Equilibria with Adverse Selection,” *International Economic Review*, 55:399-422.
- [31] Picard P. (2014): “Participating Insurance Contracts and the Rothschild-Stiglitz Equilibrium Puzzle,” *The Geneva Risk and Insurance Review*, 39:153-75.
- [32] Prescott E.C. and R.M. Townsend (1984): “Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard,” *Econometrica*, 52:21-46.
- [33] Prescott E.C. and R.M. Townsend (1984): “General Competitive Equilibria Analysis in an Economy with Private Information,” *International Economic Review*, 25:1-20.
- [34] Riley J.G. (1979): “Informational Equilibrium,” *Econometrica*, 47:331-59.
- [35] Rothschild M. and J. Stiglitz (1976): “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics*, 90:629-49.
- [36] Rustichini A. and P. Siconolfi (2008): “General Equilibrium in Economies with Adverse Selection,” *Economic Theory*, 37:1-29.
- [37] von Siemens F.A. and M. Kosfeld (2014): “Team Production in Competitive Markets with Adverse Selection,” *European Economic Review*, 68:181-98.
- [38] von Siemens F.A. and M. Kosfeld (2009): “Negative Externalities and Equilibrium Existence in Competitive Markets with Adverse Selection,” IZA Discussion Paper No. 4125.

[39] Wambach A. (2000): “Introducing Heterogeneity in the Rothschild-Stiglitz Model,” *Journal of Risk and Insurance*, 67:579-92.

[40] Wilson C. (1977): “A Model of Insurance Markets with Incomplete Information,” *Journal of Economic Theory*, 60:167-207.

Appendices

A Efficient Insurance Provision

Lemma A.1 *Let $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ be an equilibrium profile for Γ . Then,*

$$\alpha_h(\mathcal{M}(s^{1*})) \succsim_h \mathbf{a}_h^{**}, \quad h \in \{L, H\}$$

Proof. Suppose first that $\mathbf{a}_H^{**} \succ_H \alpha_H(\mathcal{M}(s^{1*}))$. Recall that the high-risk type is fully-insured at the RSW allocation: $w(a_{0H}^{**}) = w(a_{1H}^{**})$ or equivalently $u'(w(a_{0H}^{**})) = u'(w(a_{1H}^{**}))$, and thus $I_H(\mathbf{a}_H^{**}) = \frac{1-p_H}{p_H}$. For $\kappa_H < \frac{1-p_H}{p_H}$ and sufficiently small $\epsilon_H > 0$ the contract $\hat{\mathbf{a}}_H = \mathbf{a}_H^{**} + (1, \kappa_H)\epsilon_H$ gives $\mathbf{a}_H^{**} \succ_H \hat{\mathbf{a}}_H \succ_H \alpha_H(\mathcal{M}(s^{1*}))$ while $\mathbf{a}_H^{**} \succ_L \hat{\mathbf{a}}_H$ (see Lemma 1.4 in the Online Appendix). Recall also that the low-risk type is under-insured at the RSW allocation: $w(a_{0L}^{**}) > w(a_{1L}^{**})$, and thus $I_L(\mathbf{a}_L^{**}) < \frac{1-p_L}{p_L}$. Since $\mathbf{a}_L^{**} \succ_L \mathbf{a}_H^{**} \succ_L \hat{\mathbf{a}}_H$, for $\kappa_L > \frac{1-p_L}{p_L}$ and $\epsilon_L < 0$ with $|\epsilon_L|$ sufficiently small, the contract $\hat{\mathbf{a}}_L = \mathbf{a}_L^{**} + (1, \kappa_L)\epsilon_L$ gives $\mathbf{a}_L^{**} \succ_L \hat{\mathbf{a}}_L \succ_L \hat{\mathbf{a}}_H$ while $\mathbf{a}_L^{**} \succ_H \hat{\mathbf{a}}_L$ (see again Lemma 1.4 in the Online Appendix). And since $\mathbf{a}_H^{**} \sim_H \mathbf{a}_L^{**} \succ_H \hat{\mathbf{a}}_L$ while $\lim_{\epsilon_H \rightarrow 0} \hat{\mathbf{a}}_H = \mathbf{a}_H^{**}$, letting $\Delta_H = U_H(\mathbf{a}_H^{**}) - U_H(\hat{\mathbf{a}}_L)$ gives $U_H(\mathbf{a}_H^{**}) - U_H(\hat{\mathbf{a}}_H) < \Delta_H$ and thus $\hat{\mathbf{a}}_H \succ_H \hat{\mathbf{a}}_L$ for sufficiently small ϵ_H .

Suppose next that $\mathbf{a}_L^{**} \succ_L \alpha_L(\mathcal{M}(s^{1*}))$, and construct the contracts $\hat{\mathbf{a}}_L$ and $\hat{\mathbf{a}}_H$ as before. Lemma 1.4 in the Online Appendix now dictates that $\mathbf{a}_L^{**} \succ_L \hat{\mathbf{a}}_L \succ_L \alpha_L(\mathcal{M}(s^{1*}))$ and $\mathbf{a}_L^{**} \succ_H \hat{\mathbf{a}}_L$. And as $\mathbf{a}_L^{**} \succ_L \mathbf{a}_H^{**}$ while $\lim_{\epsilon_L \rightarrow 0} \hat{\mathbf{a}}_L = \mathbf{a}_L^{**}$, clearly $\hat{\mathbf{a}}_L \succ_L \mathbf{a}_H^{**}$ for sufficiently small $|\epsilon_L|$. Observe also that, since $\mathbf{a}_H^{**} \sim_H \mathbf{a}_L^{**} \succ_H \hat{\mathbf{a}}_L$, it must be $\mathbf{a}_H^{**} \succ_H \hat{\mathbf{a}}_H \succ_H \hat{\mathbf{a}}_L$ and $\mathbf{a}_H^{**} \succ_L \hat{\mathbf{a}}_H$ (recall once again Lemma 1.4 in the Online Appendix).

Putting our observations from the two cases above together, it is trivial to check that in either case the menu $\hat{m} = \{\hat{\mathbf{a}}_L, \hat{\mathbf{a}}_H\}$ is separating and strictly-preferred by at least one risk-type to $\{\alpha_L(\mathcal{M}(s^{1*})), \alpha_H(\mathcal{M}(s^{1*}))\}$. Suppose now that some firm $\hat{n} \in \mathcal{N}$ deviates to the strategy $\hat{s}_{\hat{n}}$ entailing the actions $\hat{s}_{\hat{n}|\mathcal{M}}^1 = \{\hat{m}\}$ and $\hat{s}_{\hat{n}}^2(\hat{m}, s_{-\hat{n}}^{1*}) = NW$. As $\hat{\mathbf{a}}_h \succ_h \alpha_h(\mathcal{M}(s^{1*}))$ implies that $\alpha_h(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^{1*})) = \hat{\mathbf{a}}_h$, the firm’s belief vector regarding the distribution of risk-types amongst its applicants at the subgame $\Gamma(\hat{s}_{\hat{n}}^1, s_{-\hat{n}}^{1*})$ cannot but be some $\{\lambda_{L\hat{n}}(\hat{m}), \lambda_{H\hat{n}}(\hat{m})\} \in [0, 1] \times [0, 1]$: $\lambda_{L\hat{n}}(\hat{m}) + \lambda_{H\hat{n}}(\hat{m}) = 1$. Given though that

$$\Pi_h(\hat{\mathbf{a}}_h) = \Pi_h(\mathbf{a}_h^{**}) + \left(\frac{1-p_h}{p_h} - \kappa_h \right) p_h \epsilon_h = \left(\frac{1-p_h}{p_h} - \kappa_h \right) p_h \epsilon_h > 0 \quad h \in \{L, H\}$$

this leads to the absurdity that its expected profits will be

$$\begin{aligned} \Pi_{\hat{n}}^{NW} \left(\hat{s}_{\hat{n}|\mathcal{M}}^1, s_{-\hat{n}}^*, \alpha \right) &= \frac{\lambda_{L\hat{n}}(\hat{m}) \Pi_L(\hat{\mathbf{a}}_L)}{|\mathcal{N}_L(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))|} \mathbf{1}_{\hat{m} \in \mathcal{M}_L(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))} \\ &+ \frac{\lambda_{H\hat{n}}(\hat{m}) \Pi_H(\hat{\mathbf{a}}_H)}{|\mathcal{N}_H(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))|} \mathbf{1}_{\hat{m} \in \mathcal{M}_H(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))} > 0 \end{aligned}$$

the inequality because $\hat{m} \in \mathcal{M}_h(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))$ for at least one h . ■

Lemma A.2 *Let $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ be an equilibrium profile for Γ . Then,*

$$m^* = \{\alpha_L(\mathcal{M}(s^*)), \alpha_H(\mathcal{M}(s^*))\} \in \mathcal{M}_\lambda$$

Proof. To argue ad absurdum, suppose that $\Pi(m^*|\lambda) < 0$. It must be then $\Pi_L(\alpha_L(\mathcal{M}(s^*))) < 0$ or $\Pi_H(\alpha_H(\mathcal{M}(s^*))) < 0$ or both. Let $\Pi_H(\alpha_H(\mathcal{M}(s^*))) < 0$. As obviously $\alpha_H(\mathcal{M}(s^*)) \neq \mathbf{0}$, there is at least one firm trading with the high-risk type on the equilibrium path. By free entry, however, no firm can be trading only with this type. Clearly, we must have $\emptyset \neq \mathcal{N}_H(\mathcal{M}(s^*)) \subseteq \mathcal{N}_L(\mathcal{M}(s^*))$. That is, $0 < |\mathcal{N}_H(\mathcal{M}(s^*))| \leq |\mathcal{N}_L(\mathcal{M}(s^*))|$. Suppose now that $n \in \mathcal{N}$ is trading with both types on the equilibrium path. Letting $m \in \mathcal{M}(s_n^*) \cap \mathcal{M}_L(\mathcal{M}(s^*)) \cap \mathcal{M}_H(\mathcal{M}(s^*))$, by equation (5) in the main text and free entry, it must be

$$\begin{aligned} 0 \leq \Pi_n(s^*, \alpha^*) &= \frac{\lambda \Pi_L(\mathbf{a}_L(m))}{|\mathcal{N}_L(\mathcal{M}(s^*))|} + \frac{(1-\lambda) \Pi_L(\mathbf{a}_H(m))}{|\mathcal{N}_H(\mathcal{M}(s^*))|} \\ &= \frac{\lambda \Pi_L(\alpha_L(\mathcal{M}(s^*)))}{|\mathcal{N}_L(\mathcal{M}(s^*))|} + \frac{(1-\lambda) \Pi_L(\alpha_H(\mathcal{M}(s^*)))}{|\mathcal{N}_H(\mathcal{M}(s^*))|} \\ &\leq \frac{\lambda \Pi_L(\alpha_L(\mathcal{M}(s^*))) + (1-\lambda) \Pi_H(\alpha_H(\mathcal{M}(s^*)))}{|\mathcal{N}_L(\mathcal{M}(s^*))|} = \frac{\Pi(m^*|\lambda)}{|\mathcal{N}_L(\mathcal{M}(s^*))|} \end{aligned}$$

which contradicts though our hypothesis that $m^* \notin \mathcal{M}_\lambda$. A trivially similar argument leads to the same contradiction if we let instead $\Pi_L(\alpha_L(\mathcal{M}(s^*))) < 0$. ■

Lemma A.3 *Let $m \in \mathcal{M}_\lambda \setminus \{\bar{m}\}$ be such that $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(m)$ [resp. $\mathbf{a}_L(m) \succ_L \bar{\mathbf{a}}_L$]. There exists then $\hat{m} \in \mathcal{M}_\lambda^+$ such that $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(\hat{m}) \succ_L \mathbf{a}_L(m)$ [resp. $\bar{\mathbf{a}}_H \succ_H \mathbf{a}_H(\hat{m}) \succ_H \mathbf{a}_H(m)$] and $\Pi_H(\mathbf{a}_H(\hat{m})) \leq 0$ [resp. $\Pi_L(\mathbf{a}_L(\hat{m})) < 0$].*

Moreover, the menu \hat{m} can be chosen so that $\mathbf{a}_L(\hat{m})$ [resp. $\mathbf{a}_H(\hat{m})$] is arbitrarily close to $\bar{\mathbf{a}}_L$ [resp. $\bar{\mathbf{a}}_H$], and thus $\mathbf{a}_L(\hat{m}) \succ_L \mathbf{a}_L(m')$ [resp. $\mathbf{a}_H(\hat{m}) \succ_H \mathbf{a}_H(m')$] for any given $m' \in \mathcal{M}$ with $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(m') \succ_L \mathbf{a}_L(m)$ [resp. $\bar{\mathbf{a}}_H \succ_H \mathbf{a}_H(m') \succ_H \mathbf{a}_H(m)$].

Proof.

Step 1. Suppose first that $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(m)$. Letting $\{\mathbf{a}_L^\mu, \mathbf{a}_H^\mu\}$ be the IE(μ)-optimum for some $\mu \in (\lambda, \bar{\mu})$ sufficiently close to $\bar{\mu}$, we get

$$U_L(\mathbf{a}_H^\mu) < U_L(\mathbf{a}_L^\mu) \leq U_L(\bar{\mathbf{a}}_L), \quad U_H(\mathbf{a}_L^\mu) = U_H(\mathbf{a}_H^\mu), \quad U_L(\mathbf{a}_L(m)) < U_L(\mathbf{a}_L^\mu)$$

All relations but the last inequality above follow from the choice of μ (see Claim 2.1 and Lemma 1.5 in the Online Appendix; the weak inequality accounts for the case in which the MWS allocation

coincides with the RSW one). To verify the last inequality, let $\Delta_1 = U_L(\bar{\mathbf{a}}_L) - U_L(\mathbf{a}_L(m))$ and observe that the value $V : [0, 1] \mapsto \mathbb{R}$ and solution $\mathbf{a}_{h \in \{L, H\}}^* : [0, 1] \mapsto \mathcal{A}$ of the $\text{IE}(\mu)$ -problem are continuous functions at any $\mu \in (\lambda, 1)$ (see Lemma 2.2 in the Online Appendix) whereas $U_{h \in \{L, H\}}(\cdot)$ is continuous everywhere. Observe also that $U_L(\mathbf{a}_L^\mu) \neq U_H(\mathbf{a}_H^\mu)$; otherwise, we would have $U_L(\mathbf{a}_L^\mu) = U_H(\mathbf{a}_H^\mu) = U_H(\mathbf{a}_L^\mu)$ - i.e., $(p_H - p_L)[u(w_{0L}^\mu) - u(w_{1L}^\mu)]$ which implies that \mathbf{a}_L^μ provides full insurance, an absurdity given that $\mu \neq \lambda$. Taking therefore μ sufficiently close to $\bar{\mu}$ guarantees that

$$\max \left\{ (1 - \mu) |U_L(\mathbf{a}_L^\mu) - U_H(\mathbf{a}_H^\mu)|, |V(\bar{\mu}) - V(\mu)| \right\} < \Delta_1/2$$

and thus

$$\begin{aligned} U_L(\mathbf{a}_L^\mu) &> \mu U_L(\mathbf{a}_L^\mu) + (1 - \mu) U_H(\mathbf{a}_H^\mu) - \frac{\Delta_1}{2} \\ &= V(\mu) - \frac{\Delta_1}{2} > V(\bar{\mu}) - \Delta_1 = U_L(\bar{\mathbf{a}}_L) - \Delta_1 = U_L(\mathbf{a}_L(m)) \end{aligned}$$

Consider now the contract $\hat{\mathbf{a}}_L = \mathbf{a}_L^\mu + (1, \kappa)\epsilon$ for some $\kappa \in (0, I_H(\mathbf{a}_L^\mu))$ and $\epsilon > 0$. Since $\mathbf{a}_L^\mu \succ_L \mathbf{a}_L(m)$ a sufficiently small ϵ gives $\mathbf{a}_L^\mu \succ_L \hat{\mathbf{a}}_L \succ_L \mathbf{a}_L(m)$ and $\mathbf{a}_L^\mu \succ_H \hat{\mathbf{a}}_L$ (see Lemma 1.4 in the Online Appendix). And since also $\mathbf{a}_H^\mu \sim_H \mathbf{a}_L^\mu$, the menu $\hat{m} = \{\hat{\mathbf{a}}_L, \mathbf{a}_H^\mu\}$ satisfies (IC_H) . It also satisfies $(\text{IR}_{h \in \{L, H\}})$ for μ sufficiently close to $\bar{\mu}$ (see Lemma 2.5(i) and Remark IV(i) in the Online Appendix). To see that it satisfies (IC_L) , recall the proof of Lemma 1.3 in the Online Appendix. To establish that $\hat{\mathbf{a}}_L \succ_L \mathbf{a}_L$ we used that $U_L(\hat{\mathbf{a}}_L) - U_L(\mathbf{a}_L^\mu) \in (-\Delta_2, 0)$ for sufficiently small ϵ and where $\Delta_2 = U_L(\mathbf{a}_L^\mu) - U_L(\mathbf{a}_L)$. It is trivial to check that the argument remains valid if we define instead $\Delta_2 = U_L(\mathbf{a}_L^\mu) - \max\{U_L(\mathbf{a}_H^\mu), U_L(\mathbf{a}_L(m))\}$ to obtain $U_L(\hat{\mathbf{a}}_L) > \max\{U_L(\mathbf{a}_H^\mu), U_L(\mathbf{a}_L(m))\}$, again for sufficiently small ϵ . To show next that $\Pi(\hat{m}|\lambda) > 0$ notice that, compared to $\{\mathbf{a}_L^\mu, \mathbf{a}_H^\mu\}$ which breaks even across the two types, the new menu faces exactly the same expected losses from the high-risk type. It expects though even higher profits from the low-risk since

$$\begin{aligned} \Pi_L(\hat{\mathbf{a}}_L) - \Pi_L(\mathbf{a}_L^\mu) &= \left(\frac{1 - p_L}{p_L} - \kappa \right) p_L \epsilon > \left[\frac{1 - p_L}{p_L} - I_L(\mathbf{a}_L^\mu) \right] p_L \epsilon \\ &= (1 - p_L) \left[1 - \frac{u'(w_{0L}^\mu)}{u'(w_{1L}^\mu)} \right] \epsilon > 0 \end{aligned}$$

The first inequality here is because $\kappa < I_H(\mathbf{a}_L^\mu) < I_L(\mathbf{a}_L^\mu)$. The second follows from risk-aversion and the fact that, $\{\mathbf{a}_L^\mu, \mathbf{a}_H^\mu\}$ being $\text{IE}(\mu)$ with $\mu > \lambda$, it must be $w_{0L}^\mu > w_{1L}^\mu$ (see Claim 2.1 in the Online Appendix). Finally, we have $\mathbf{a}_h(\hat{m}) \in \mathbb{R}_{++}^2$ for either h (see Lemma 2.5(v) in the Online Appendix) as well as $\Pi_H(\mathbf{a}_H^\mu) \leq 0$ (see Lemma 2.4 in the Online Appendix).

Step 2. Suppose next that $\mathbf{a}_L(m) \succ_L \bar{\mathbf{a}}_L$ with $\mathbf{a}_L(m) \neq \bar{\mathbf{a}}_L$. It cannot but be $\bar{\mu} < 1$ and $U_H(\mathbf{a}_H(m)) < U_H(\bar{\mathbf{a}}_H)$. Let then $\Delta_3 = U_H(\bar{\mathbf{a}}_H) - U_H(\mathbf{a}_H(m))$. Taking $\mu \in (\bar{\mu}, 1)$ sufficiently close to $\bar{\mu}$ ensures that $\Pi_L(\mathbf{a}_L^\mu) < 0 < \Pi_H(\mathbf{a}_H^\mu)$ (see Lemma 2.4 in the Online Appendix), while $0 < U_H(\bar{\mathbf{a}}_H) - U_H(\mathbf{a}_H^\mu) < \Delta_3$ (see Lemma 1.5 in the Online Appendix) and thus $U_H(\mathbf{a}_H(m)) < U_H(\mathbf{a}_H^\mu)$. Notice also that $\mathbf{a}_L^\mu \succ_L \mathbf{a}_H^\mu$ (see Claim 2.1 in the Online Appendix) and let $\Delta_4 = U_L(\mathbf{a}_L^\mu) - U_L(\mathbf{a}_H^\mu)$. Consider now the menu $\hat{m} = \{\hat{\mathbf{a}}_L, \mathbf{a}_H^\mu\}$ with $\hat{\mathbf{a}}_L = \mathbf{a}_L^\mu + (1, \kappa)\epsilon$

for some $\kappa < I_H(\mathbf{a}_L^\mu)$ and $\epsilon > 0$. A sufficiently small $|\epsilon|$ gives $\mathbf{a}_L^\mu \succ_L \hat{\mathbf{a}}_L \succ_L \mathbf{a}_H^\mu$ and $\mathbf{a}_L^\mu \succ_H \hat{\mathbf{a}}_L$ (Lemma 1.4 in the Online Appendix), as well as $U_L(\mathbf{a}_L^\mu) - U_L(\hat{\mathbf{a}}_L) < \Delta_4$ and thus $\hat{\mathbf{a}}_L \succ_L \mathbf{a}_H^\mu$. Since also $\mathbf{a}_H^\mu \sim_H \mathbf{a}_L^\mu$, the menu \hat{m} is separating. It meets moreover either (IR_h) constraint (see Lemma 2.5(ii) and Remark IV(ii) in the Online Appendix) and gives $\mathbf{a}_h(\hat{m}) \in \mathbb{R}_{++}^2$ for either h (see Remark IV(iii) in the Online Appendix). Finally, since $\bar{\mu} > \lambda$, the same argument as in the previous case shows that $\Pi(\hat{m}|\lambda) > 0$; in the present case, since $\Pi_L(\mathbf{a}_L^\mu) < 0$ (see again Lemma 2.4 in the Online Appendix), that $\Pi_L(\mathbf{a}_L(\hat{m})) < 0$ follows from taking $|\epsilon|$ sufficiently small.

Step 3. We have established the validity of the claim on the space $\mathcal{M}_\lambda \setminus \{m \in \mathcal{M}_\lambda : \mathbf{a}_L(m) \neq \bar{\mathbf{a}}_L\}$. To complete the argument for the first part of the claim, let $m = \{\bar{\mathbf{a}}_L, \mathbf{a}_H\}$ with $\mathbf{a}_H \in \mathcal{A} \setminus \{\bar{\mathbf{a}}_H\}$. The MWS menu being the unique $\text{IE}(\bar{\mu})$ -optimum, we must have $\bar{\mathbf{a}}_H \succ_H \mathbf{a}_H(m)$ and thus $\bar{\mathbf{a}}_L \sim_H \bar{\mathbf{a}}_H \succ_H \mathbf{a}_H(m)$. But then m violates (IC_H), a contradiction of the hypothesis that $m \in \mathcal{M}_\lambda \subset \mathcal{M}$. Step 4. Recall that, in either of Steps 1-2 above, the construction of \hat{m} allows μ and ϵ to be taken arbitrarily close to $\bar{\mu}$ and zero, respectively. In other words, \hat{m} can be chosen arbitrarily close to \bar{m} . The second part of the claim follows from the continuity of the functions $U_{h \in \{L, H\}}(\cdot)$. ■

Proof of Claim 3.1

Let $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ be an equilibrium profile for Γ . To argue ad absurdum, let suppose first that $\bar{\mathbf{a}}_L \succ_L \alpha_L(\mathcal{M}(s^{1*}))$. Since $\{\alpha_L(\mathcal{M}(s^*)), \alpha_H(\mathcal{M}(s^*))\} \in \mathcal{M}_\lambda$ (recall Lemma A.2), by Lemma A.3 there exists $\hat{m} = (\hat{\mathbf{a}}_L, \hat{\mathbf{a}}_H)$ which is incentive-compatible, individually-rational for either risk-type, and such that $\hat{\mathbf{a}}_L \succ_L \alpha_L(\mathcal{M}(s^{1*}))$ while $\Pi_H(\hat{\mathbf{a}}_H) \leq 0 < \Pi(\hat{m}|\lambda)$. Suppose then that firm $n \in \mathcal{N}$ deviates to the strategy \hat{s}_n which entails the action $\hat{s}_{n|\mathcal{M}}^1 = \{\hat{m}\} = \hat{s}_{n|C}^1$ at stage 1 (and subsequently the choice $\hat{s}_n^2(\hat{m}, \cdot) = \text{NW}$ at stage 3 everywhere on S_{-n}^1). Since $\hat{\mathbf{a}}_L \succ_L \alpha_L(\mathcal{M}(s^{1*}))$ ensures that $\alpha_L(\mathcal{M}(\hat{s}_n, s_{-n}^*)) = \hat{\mathbf{a}}_L$, the deviant should expect profits

$$\Pi_n^C(\{\hat{m}\}, s_{-n}^*, \alpha) = \lambda \Pi_L(\hat{\mathbf{a}}_L) + \frac{(1-\lambda) \Pi_H(\hat{\mathbf{a}}_H)}{|\mathcal{N}_H(\mathcal{M}(\hat{s}_n, s_{-n}^*))|} \times \mathbf{1}_{\hat{m} \in \mathcal{M}_H(\mathcal{M}(\hat{s}_n, s_{-n}^*))}$$

and we need to account for two possible cases. If $\hat{m} \notin \mathcal{M}_H(\mathcal{M}(\hat{s}_n, s_{-n}^*))$, we have

$$\Pi_n^C(\{\hat{m}\}, s_{-n}^*, \alpha) = \lambda \Pi_L(\hat{\mathbf{a}}_L) \geq \lambda \Pi_L(\hat{\mathbf{a}}_L) + (1-\lambda) \Pi_H(\hat{\mathbf{a}}_H) = \Pi(\hat{m}|\lambda) > 0$$

Otherwise, $|\mathcal{N}_H(\mathcal{M}(\hat{s}_n, s_{-n}^*))| \geq 1$ and thus

$$\begin{aligned} \Pi_n^C(\{\hat{m}\}, s_{-n}^*, \alpha) &= \lambda \Pi_L(\hat{\mathbf{a}}_L) + \frac{(1-\lambda) \Pi_H(\hat{\mathbf{a}}_H)}{|\mathcal{N}_H(\mathcal{M}(\hat{s}_n, s_{-n}^*))|} \\ &\geq \lambda \Pi_L(\hat{\mathbf{a}}_L) + (1-\lambda) \Pi_H(\hat{\mathbf{a}}_H) = \Pi(\hat{m}|\lambda) > 0 \end{aligned}$$

In either case therefore the deviation is profitable. To reach moreover the same conclusion when $\alpha_L(\mathcal{M}(s^{1*})) \succsim_L \bar{\mathbf{a}}_L$ with $\alpha_L(\mathcal{M}(s^{1*})) \neq \bar{\mathbf{a}}_L$ requires but a trivially similar argument. ■

Corollary A.1 $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ is an equilibrium profile for Γ only if

$$\{\alpha_L(\mathcal{M}(s^*)), \alpha_H(\mathcal{M}(s^*))\} \in \mathcal{M}_\lambda^* := \{m \in \mathcal{M}_\lambda : \bar{\mathbf{a}}_L \succsim_L \alpha_L(m)\}$$

Proof. As $\mathcal{M}(s^*) \subseteq \mathcal{M}(s^{1*})$ by the definition of our mechanism, the result follows from Claim 3.1 and Lemma A.2. ■

Proof of Proposition 3.1

To support the sequential equilibrium scenario in question, we will consider some $\hat{n} \in \mathcal{N}$ deviating with the strategy $\hat{s}_{\hat{n}}$ that entails offering some the set of menus $M \in \overline{\mathcal{M}}$ at stage 1 (i.e., $\hat{s}_{\hat{n}}^1|_{\mathcal{M}} = \hat{M}$) and honoring $m_L, m_H \in \hat{M}$ (i.e., $\hat{s}_{\hat{n}}^2(m_L, s_{-\hat{n}}^{1*}) = NW = \hat{s}_{\hat{n}}^2(m_H, s_{-\hat{n}}^{1*})$) where $(m_L, m_H) \in \mathcal{M}_L(\mathcal{M}(\hat{s}_{\hat{n}})) \times \mathcal{M}_H(\mathcal{M}(\hat{s}_{\hat{n}}))$. This implies of course that $\hat{m} = \{\mathbf{a}_L(m_L), \mathbf{a}_H(m_H)\}$ satisfies (IC_h) for either h . We will let the trembles in (7) be given by the sequences $\{\epsilon'_{kh}\}_{(k,h) \in \mathbb{N} \times \{L,H\}}$ and $\{\epsilon''_{kh}\}_{(k,h) \in \mathbb{N} \times \{L,H\}}$ for $\hat{m} = m_L$ and $\hat{m} = m_H$, respectively (recall footnote 19 in the main text). As either risk-type chooses each of \bar{m} , m_L , and m_H with positive probability, the payoffs are given by

$$\begin{aligned} \Pi_{\hat{n}k}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) &= \frac{\lambda_{L\hat{n}}^k(m_L) \Pi_L(\mathbf{a}_L(m_L))}{|\mathcal{N}_L(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))|} + \frac{\lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H))}{|\mathcal{N}_H(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))|} \\ \Pi_{n^*k}^{NW}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha^*) &= \frac{\lambda_{Ln^*}^k(\bar{m}) \Pi_L(\bar{\mathbf{a}}_L)}{|\mathcal{N}_L(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))|} + \frac{\lambda_{Hn^*}^k(\bar{m}) \Pi_H(\bar{\mathbf{a}}_H)}{|\mathcal{N}_H(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*))|} \end{aligned}$$

where, for either h ,

$$\mathcal{N}_h(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*)) = \begin{cases} \{n\} & \text{if } \alpha_h(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*)) = \mathbf{a}_h(m_h) \neq \bar{\mathbf{a}}_h \\ \mathcal{N} & \text{if } \alpha_h(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*)) = \mathbf{a}_h(m_h) = \bar{\mathbf{a}}_h \\ \mathcal{N} \setminus \{n\} & \text{if } \alpha_h(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*)) = \bar{\mathbf{a}}_h \neq \mathbf{a}_h(m_h) \end{cases}$$

Step 1. Let the deviation attract only the high-risk type away from \bar{m} : $\bar{\mathbf{a}}_L \succ_L \mathbf{a}_L(m_L)$ and $\mathbf{a}_H(m_H) \succ_H \bar{\mathbf{a}}_H$. Since $\mathbf{a}_H(m_H) \succ_H \mathbf{a}_H^*$, it cannot but be $\Pi_H(\mathbf{a}_H(m_H)) < 0$; the deviation is potentially profitable only if $\Pi_L(\mathbf{a}_L(m_L)) > 0$. And $\Pi_L(\mathbf{a}_L(m_L)) > 0 > \Pi_H(\mathbf{a}_H(m_H))$ allows us to define the quantity $0 < \Delta = -(1 - \lambda) \Pi_H(\mathbf{a}_H(m_H)) / [\lambda \Pi_L(\mathbf{a}_L(m_L))]$.

We will support the scenario that has $\{m_L, m_H\}$ withdrawn and \bar{m} honoured, while either risk-type applies for \bar{m} (i.e., $\sigma_h(\bar{\mathbf{a}}_h) = 1 = 1 - \sigma_h(\mathbf{a}_h(m_h))$) for either h).

(i). If m_H is offered without commitment, let $\epsilon'_{kL}/\epsilon'_{kH} \leq \epsilon''_{kL}/\epsilon''_{kH} < \Delta$ and $\epsilon'_{kL} + \epsilon''_{kL} \leq \epsilon'_{kH} + \epsilon''_{kH}$. The first two inequalities ensure that $\lambda_{L\hat{n}}^k(m_L) \leq \lambda_{L\hat{n}}^k(m_H)$ and $\lambda_{L\hat{n}}^k(m_L)/\lambda_{H\hat{n}}^k(m_H) < -\Pi_H(\mathbf{a}_H(m_H))/\Pi_L(\mathbf{a}_L(m_L))$, respectively, everywhere along the sequences.³⁶ As a result,

$$\begin{aligned} \Pi_{\hat{n}k}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) &= \lambda_{L\hat{n}}^k(m_L) \Pi_L(\mathbf{a}_L(m_L)) + \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) \\ &\leq \lambda_{L\hat{n}}^k(m_H) \Pi_L(\mathbf{a}_L(m_L)) + \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) \\ &= \lambda_{H\hat{n}}^k(m_H) \left[\frac{\lambda_{L\hat{n}}^k(m_H)}{\lambda_{H\hat{n}}^k(m_H)} \Pi_L(\mathbf{a}_L(m_L)) + \Pi_H(\mathbf{a}_H(m_H)) \right] < 0 \end{aligned}$$

By contrast, $\epsilon'_{kL} + \epsilon''_{kL} < \epsilon'_{kH} + \epsilon''_{kH}$ means that $\lambda_{n^*}^k = (1 - \epsilon'_{kL} - \epsilon''_{kL}) / (1 - \epsilon'_{kH} - \epsilon''_{kH}) > 1$ and thus $\lambda_{Ln^*}^k(\bar{m}) > \lambda$ and $\lambda_{Hn^*}^k(\bar{m}) < 1 - \lambda$ along the sequence. And since $\Pi_L(\bar{\mathbf{a}}_L) \geq 0 \geq \Pi_H(\bar{\mathbf{a}}_H)$ (see Lemma 2.4 in the Online Appendix), we get

$$\begin{aligned} \Pi_{n^*k}^{NW}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha^*) &= (N - 1)^{-1} \left[\lambda_{Ln^*}^k(\bar{m}) \Pi_L(\bar{\mathbf{a}}_L) + \lambda_{Hn^*}^k(\bar{m}) \Pi_H(\bar{\mathbf{a}}_H) \right] \\ &\geq (N - 1)^{-1} [\lambda \Pi_L(\bar{\mathbf{a}}_L) + (1 - \lambda) \Pi_H(\bar{\mathbf{a}}_H)] = (N - 1)^{-1} \Pi(\bar{m}|\lambda) = 0 \end{aligned}$$

³⁶Observe that $\lambda_{L\hat{n}}^k(m)$ in (7) is strictly increasing with respect to $\lambda_{\hat{n}}^k$.

(ii). Let now m_H be offered with commitment. Under the equilibrium scenario in consideration we have $\mathbf{a}_H(m_H) \succ_H \bar{\mathbf{a}}_H = \alpha_H(\mathcal{M}(s_{-\hat{n}}^*))$. That is, $\alpha_H(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*)) = \mathbf{a}_H(m_H) \neq \bar{\mathbf{a}}_H$ and hence the only relevant sequence of trembles is $\{\epsilon'_{kh}\}_{(k,h) \in \mathbb{N} \times \{L,H\}}$. Letting $\lim_{k \rightarrow \infty} \epsilon'_{kL}/\epsilon'_{kH} = 0$ ensures that also $\lim_{k \rightarrow \infty} \lambda_{L\hat{n}}^k(m_L) = 0$. That is, $\lambda_{L\hat{n}}^k(m_L) < \lambda\Delta$ (at least along a subsequence) and thus

$$\Pi_{\hat{n}k}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) = \lambda_{L\hat{n}}^k(m_L) \Pi_L(\mathbf{a}_L(m_L)) + (1 - \lambda) \Pi_H(\mathbf{a}_H(m_H)) < 0$$

Moreover, $\epsilon'_{kL} < \epsilon'_{kH}$ (at least along a subsequence) and thus

$$\Pi_{n^*k}^{NW}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha^*) = (N - 1)^{-1} \lambda_{Ln^*}^k(\bar{m}) \Pi_L(\bar{\mathbf{a}}_L) > (N - 1)^{-1} \lambda \Pi_L(\bar{\mathbf{a}}_L) \geq 0$$

Step 2. Let now the deviation attract at least the low-risk type away from \bar{m} : $\mathbf{a}_L(m_L) \succ_L \bar{\mathbf{a}}_L$. We will support the scenario that has $\{m_L, m_H\}$ withdrawn, \bar{m} withdrawn if only if $m_L \in \hat{s}_{\hat{n}|C}^1$ and $\bar{m} \neq m^{**}$, while either risk-type applies for \bar{m} .

(i). If $\mathbf{a}_H^{**} \succ_H \mathbf{a}_H(m_H)$, we have $\mathbf{a}_L(m_L) \succ_L \bar{\mathbf{a}}_L \succ_L \bar{\mathbf{a}}_H \succ_H \mathbf{a}_H^{**}$ and $\mathbf{a}_H^{**} \succ_H \mathbf{a}_H(m_H) \succ_L \mathbf{a}_L(m_L)$. The menu $\{\mathbf{a}_L(m_L), \mathbf{a}_H^{**}\}$ is incentive-compatible for either risk type and, the RSW allocation being optimal for the RSW problem, it must be $\Pi_L(\mathbf{a}_L(m_L)) < 0$. Clearly, $\mathbf{a}_L(m_L) \neq \bar{\mathbf{a}}_L$ while the deviation is not potentially profitable unless $\Pi_H(\mathbf{a}_H(m_H)) > 0$.

(i.a). Let $m_L \in \hat{s}_{\hat{n}|NC}^1$.

(i.a.i). If $\bar{\mu} = 1$, it cannot but be $\Pi(\hat{m}|\lambda) < 0$; more precisely, $\mathcal{P}(\hat{m}) \subset [0, \lambda)$. Let then $\epsilon'_{kL}/\epsilon'_{kH} = 1 = \epsilon''_{kL}/\epsilon''_{kH}$. Since $\lambda_{L\hat{n}}^k(m_L) = \lambda = \lambda_{L\hat{n}}^k(m_H)$, we have

$$\begin{aligned} \Pi_{\hat{n}k}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) &= \lambda_{L\hat{n}}^k(m_L) \Pi_L(\mathbf{a}_L(m_L)) + \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) \\ &= \lambda \Pi_L(\mathbf{a}_L(m_L)) + (1 - \lambda) \Pi_H(\mathbf{a}_H(m_H)) = \Pi(\hat{m}|\lambda) < 0 \end{aligned}$$

By contrast, $\epsilon'_{kL} + \epsilon''_{kL} = \epsilon'_{kH} + \epsilon''_{kH}$ means that $\lambda_{n^*}^k = 1$. That is, $\lambda_{Ln^*}^k(\bar{m}) = \lambda$ and

$$\Pi_{n^*k}^{NW}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha^*) = (N - 1)^{-1} \left[\lambda_{Ln^*}^k(\bar{m}) \Pi_L(\bar{\mathbf{a}}_L) + \lambda_{Hn^*}^k(\bar{m}) \Pi_H(\bar{\mathbf{a}}_H) \right] = (N - 1)^{-1} \Pi(\bar{m}|\lambda) = 0$$

It is optimal for $n^* \in \mathcal{N}$ to honor \bar{m} at stage 3.

(i.a.ii). If $\bar{\mu} < 1$, then $\bar{m} = m^{**}$; any firm offering this menu is always indifferent between honoring it or not at stage 3 (recall the remark that follows the definition of $\mathcal{P}(\cdot)$ in the main text: in the present case, $\mathcal{P}(\bar{m}) = [0, 1]$). Let now $\epsilon'_{kL}/\epsilon'_{kH} \geq \epsilon''_{kL}/\epsilon''_{kH} > \Delta$. Since $\Pi_L(\mathbf{a}_L(m_L)) < 0$, adapting trivially the argument in Step 1(i) shows that $\Pi_{\hat{n}k}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) < 0$ (see the left panel of Figure 5).

(i.b). Let $m_L \in \hat{s}_{\hat{n}|C}^1$, and recall the argument in Step 1(ii). Here we have $\alpha_L(\mathcal{M}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*)) = \mathbf{a}_L(m_L) \neq \bar{\mathbf{a}}_L$. The relevant sequence now is $\{\epsilon''_{kh}\}_{(k,h) \in \mathbb{N} \times \{L,H\}}$. Letting $\lim_{k \rightarrow \infty} \epsilon''_{kL}/\epsilon''_{kH} = \infty$ means that $\lim_{k \rightarrow \infty} \lambda_{H\hat{n}}^k(m_H) = 0$ and thus $\lambda_{H\hat{n}}^k(m_H) < (1 - \lambda)/\Delta$ (at least along a subsequence). This gives

$$\Pi_{\hat{n}k}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) = \lambda \Pi_L(\mathbf{a}_L(m_L)) + \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) < 0$$

Moreover, $\Pi_{n^*k}^{NW}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha^*) = \lambda_{Hn^*}^k(\bar{m}) \Pi_H(\bar{\mathbf{a}}_H) / [N - 1 + \mathbf{1}_{\mathbf{a}_H(m_H) = \bar{\mathbf{a}}_H}] \leq 0$ (the inequality strict unless $\bar{m} = m^{**}$) means that it is optimal for $n^* \in \mathcal{N}$ to withdraw \bar{m} at stage 3 if $\bar{\mu} = 1$ and $\bar{m} \neq m^{**}$. When $\bar{m} = m^{**}$, which embeds the case $\bar{\mu} < 1$, the firm is indifferent between honoring

or withdrawing \bar{m} .

(ii). If $\mathbf{a}_H(m_H) \succ_H \mathbf{a}_H^{**}$, we have $\Pi_H(\mathbf{a}_H(m_H)) \leq 0$; the deviation is not potentially profitable unless $\Pi_L(\mathbf{a}_L(m_L)) > 0$. In fact, it must be $\Pi_H(\mathbf{a}_H(m_H)) < 0$ because $\Pi_H(\mathbf{a}_H(m_H)) = 0$ necessitates that $\mathbf{a}_H(m_H) = \mathbf{a}_H^{**}$, an absurdity since $\mathbf{a}_L(m_L) \succ_L \bar{\mathbf{a}}_L \succ_L \mathbf{a}_L^{**}$ (for any $\{\mathbf{a}, \mathbf{a}_H^{**}\} \in \mathcal{M}$, it cannot be $\Pi_L(\mathbf{a}) \geq 0$ and $\mathbf{a} \succ_L \mathbf{a}_L^{**}$ unless $\mathbf{a} = \mathbf{a}_L^{**}$). Clearly, $\mathbf{a}_H(m_H) \neq \mathbf{a}_H^{**}$ and $\Pi_L(\mathbf{a}_L(m_L)) > 0 > \Pi_H(\mathbf{a}_H(m_H))$. The argument proceeds as in Step 1(i) if $m_L \in \hat{s}_{\hat{n}|NC}^1$ (see the right panels of Figures 4,5,6). If $m_L \in \hat{s}_{\hat{n}|C}^1$, let $\epsilon''_{kL}/\epsilon''_{kH} = 1$. This gives $\lambda_{H\hat{n}}^k(m_H) = 1 - \lambda$ and thus

$$\begin{aligned} \Pi_{\hat{n}k}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) &= \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) + \lambda \Pi_L(\mathbf{a}_L(m_L)) / [1 + (N-1) \mathbf{1}_{\mathbf{a}_L(m_L)=\bar{\mathbf{a}}_L}] \\ &< \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) + \lambda \Pi_L(\mathbf{a}_L(m_L)) \\ &= (1 - \lambda) \Pi_H(\mathbf{a}_H(m_H)) + \lambda \Pi_L(\mathbf{a}_L(m_L)) = \Pi(\hat{m}|\lambda) < 0 \end{aligned}$$

the last inequality being strict (more precisely, $\mathcal{P}(\hat{m}) \subset (\lambda, 1]$; see the left panel of Figure 4) unless $\mathbf{a}_h(m_h) = \bar{\mathbf{a}}_h$ for either h . With respect to \bar{m} , the argument remains the same as in Step 2(i.b). ■

Lemma A.4 *Let $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ be an equilibrium profile for Γ . Suppose also that $\exists h \in \{L, H\}$: $\alpha_h(\mathcal{M}(s^*)) \succ_h \mathbf{a}_h^{**}$. We have then $\mathcal{N}_L(\mathcal{M}(s^*)) \subseteq \mathcal{N}_H(\mathcal{M}(s^*)) \neq \emptyset$, with $\mathcal{N}_L(\mathcal{M}(s^*)) = \emptyset$ only if $\alpha_H(\mathcal{M}(s^*)) = \mathbf{a}_H^{**}$.*

Proof. Recall that $\mathbf{a}_h^{**} \succ_h \mathbf{0}$ for either h (see Lemma 2.5(i) in the Online Appendix). By hypothesis, therefore, firms are trading with at least one type on the equilibrium path: $\mathcal{N}_h(\mathcal{M}(s^*)) \neq \emptyset$ for some $h \in \{L, H\}$. In fact, it must be $\mathcal{N}_H(\mathcal{M}(s^*)) \neq \emptyset$. For otherwise, we would have $\mathcal{N}_L(\mathcal{M}(s^*)) \neq \emptyset = \mathcal{N}_H(\mathcal{M}(s^*))$. And since there is no trade for the high-risk type (equivalently, $\alpha_H(\mathcal{M}(s^*)) = \mathbf{0}$), it must be $\alpha_L(\mathcal{M}(s^*)) \succ_L \mathbf{a}_L^{**} \sim_H \mathbf{a}_H^{**} \succ_H \mathbf{0}$. Which contradicts, though, that there is no trade with the high-risk type; for she strictly prefers the contract $\alpha_L(\mathcal{M}(s^*))$ to her endowment point. ■

Lemma A.5 *Let $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ be an equilibrium profile for Γ . If $\alpha_H(\mathcal{M}(s^*)) \succ_H \mathbf{a}_H^{**}$ then $\emptyset \neq \mathcal{N}_L(\mathcal{M}(s^*)) = \mathcal{N}_H(\mathcal{M}(s^*))$.*

Proof. By the preceding lemma, it cannot but be $\emptyset \neq \mathcal{N}_L(\mathcal{M}(s^*)) \subseteq \mathcal{N}_H(\mathcal{M}(s^*))$. Recall also that $\mathbf{a}_H^{**} = \arg \max_{\mathbf{a} \in \mathcal{A}: \Pi_H(\mathbf{a}) \geq 0} U_H(\mathbf{a})$ with $\Pi_H(\mathbf{a}_H^{**}) = 0$. Clearly, $\alpha_H(\mathcal{M}(s^*)) \succ_H \mathbf{a}_H^{**}$ implies that $\Pi_H(\mathbf{a}_H(m)) < 0$ for any $m \in \mathcal{M}_H(\mathcal{M}(s^*))$. Hence, by free entry and equation (5) in the main text, no firm can be trading with only the high-risk type on the equilibrium path. That is, $\mathcal{N}_H(\mathcal{M}(s^*)) \subseteq \mathcal{N}_L(\mathcal{M}(s^*))$ and the claim follows. ■

B Exit in Bertrand Competition

Lemma B.1 *Let $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ be an equilibrium profile for Γ such that $\alpha_H(\mathcal{M}(s^*)) \succ \mathbf{a}_H^{**}$. Then,*

$$\Pi_H(\alpha_H(\mathcal{M}(s^*))) \leq 0 \leq \Pi_L(\alpha_L(\mathcal{M}(s^*)))$$

Proof. Since $\mathbf{a}_H^{**} = \arg \max_{\mathbf{a} \in \mathcal{A}: \Pi_H(\mathbf{a}) \geq 0} U_H(\mathbf{a})$ with $\Pi_H(\mathbf{a}_H^{**}) = 0$, if $\alpha_H(\mathcal{M}(s^*)) \succsim \mathbf{a}_H^{**}$ then it cannot but be $\Pi_H(\alpha_H(\mathcal{M}(s^*))) \leq 0$. The claim holds trivially if no firm is trading with the low-risk type on the equilibrium path: $\mathcal{N}_L(\mathcal{M}(s^*)) = \emptyset$ implies that $\alpha_L(\mathcal{M}(s^*)) = \mathbf{0}$ and thus $\Pi_L(\alpha_L(\mathcal{M}(s^*))) = 0$. If $\mathcal{N}_L(\mathcal{M}(s^*)) \neq \emptyset$, on the other hand, taking $n \in \mathcal{N}_L(\mathcal{M}(s^*))$ and $m \in \mathcal{M}(s_n^*) \cap \mathcal{M}_L(\mathcal{M}(s^*))$, equation (5) in the main text gives

$$\begin{aligned} 0 \leq \Pi_n(s^*, \alpha^*) &= \frac{\lambda \Pi_L(\mathbf{a}_L(m))}{|\mathcal{N}_L(\mathcal{M}(s^*))|} + \frac{(1-\lambda) \Pi_H(\mathbf{a}_H(m))}{|\mathcal{N}_H(\mathcal{M}(s^*))|} \times \mathbf{1}_{\mathcal{M}(s_n^*) \cap \mathcal{M}_H(\mathcal{M}(s^*)) \neq \emptyset} \\ &= \frac{\lambda \Pi_L(\alpha_L(\mathcal{M}(s^*)))}{|\mathcal{N}_L(\mathcal{M}(s^*))|} + \frac{(1-\lambda) \Pi_H(\alpha_H(\mathcal{M}(s^*)))}{|\mathcal{N}_H(\mathcal{M}(s^*))|} \times \mathbf{1}_{\mathcal{M}(s_n^*) \cap \mathcal{M}_H(\mathcal{M}(s^*)) \neq \emptyset} \\ &\leq \frac{\lambda \Pi_L(\alpha_L(\mathcal{M}(s^*)))}{|\mathcal{N}_L(\mathcal{M}(s^*))|} \end{aligned}$$

where the first inequality above is due to free entry. ■

Proof of Claim 4.1

Let $m \in \mathcal{M}'$. In what follows, we will support the scenario in which the deviation [resp. m] gets withdrawn [resp. honored] at stage 3, and either risk-type chooses m .

Step 1. If $\mathbf{a}_H(m_H) \succ_H \mathbf{a}_H(m)$ and $\mathbf{a}_L(m) \succ_L \mathbf{a}_L(m_L)$, the argument proceeds as in Step 1(i) of the proof for Proposition 3.1.

Step 2. Let then $\mathbf{a}_L(m_L) \succsim_L \mathbf{a}_L(m)$. It cannot be $\mathcal{P}(\hat{m}) = [0, 1]$. For this implies $\Pi_H(\mathbf{a}_H(m_H)) = \Pi_L(\mathbf{a}_L(m_L)) > 0$ - the equality by the definition of $\mathcal{P}(\cdot)$ in the main text, and the inequality for the deviation to be profitable. Yet any solution to the problem $\max_{m \in \mathcal{M}: \Pi_h(\mathbf{a}_h(m)) \geq 0} U_L(\mathbf{a}_L(\tilde{m}))$ gives $\mathbf{a}_L(\tilde{m}) = \mathbf{a}_L^{**}$ while $\Pi_L(\mathbf{a}_L^{**}) = 0$. Clearly, $\mathbf{a}_L(\hat{m}) \succsim_L \mathbf{a}_L(m) \succsim_L \mathbf{a}_L^{**}$ renders $\Pi_H(\mathbf{a}_H(\hat{m})) = \Pi_L(\mathbf{a}_L(\hat{m})) > 0$ absurd.

Hence, either $\mathcal{P}(\hat{m}) = [p(\hat{m}), 1]$ or $\mathcal{P}(\hat{m}) = [0, p(\hat{m})]$, for some $p(\hat{m}) \in (0, 1)$.

Step 3. $\mathcal{P}(\hat{m}) = [p(\hat{m}), 1]$ necessitates that $\Pi_L(\mathbf{a}_L(m_L)) > \Pi_H(\mathbf{a}_H(m_H)) > 0$; the deviation is not potentially profitable unless $\Pi_L(\mathbf{a}_L(m_L)) > 0$. Let $\Delta = (1-\lambda)p(\hat{m}) / [\lambda(1-p(\hat{m}))]$ and choose trembles as in Step 1(i) of the proof for Proposition 3.1. Since now $\lambda_{L\hat{n}}^k(m_H) / \lambda_{H\hat{n}}^k(m_H) < p(\hat{m}) / [1-p(\hat{m})]$ everywhere along the sequences, we have

$$\begin{aligned} \Pi_{nk}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) &= \lambda_{L\hat{n}}^k(m_L) \Pi_L(\mathbf{a}_L(m_L)) + \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) \\ &\leq \lambda_{L\hat{n}}^k(m_H) \Pi_L(\mathbf{a}_L(m_L)) + \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) \\ &= \lambda_{H\hat{n}}^k(m_H) \left[\frac{\lambda_{L\hat{n}}^k(m_H)}{\lambda_{H\hat{n}}^k(m_H)} \Pi_L(\mathbf{a}_L(m_L)) + \Pi_H(\mathbf{a}_H(m_H)) \right] \\ &< \lambda_{H\hat{n}}^k(m_H) \left[\frac{p(\hat{m})}{1-p(\hat{m})} \Pi_L(\mathbf{a}_L(m_L)) + \Pi_H(\mathbf{a}_H(m_H)) \right] \\ &= \frac{\lambda_{H\hat{n}}^k(m_H)}{1-p(\hat{m})} [p(\hat{m}) \Pi_L(\mathbf{a}_L(m_L)) + (1-p(\hat{m})) \Pi_H(\mathbf{a}_H(m_H))] = 0 \end{aligned}$$

The reasoning for why it is optimal for any $n^* \in \mathcal{N}$ to honor m at stage 3 is the same as in Step 1(i) of the proof for Proposition 3.1.

Step 4. If $\mathcal{P}(\hat{m}) = [0, p(\hat{m})]$, it cannot but be $\Pi_L(\mathbf{a}_L(m_L)) < \Pi_H(\mathbf{a}_H(m_H))$; the deviation is

not potentially profitable unless $\Pi_H(\mathbf{a}_H(m_H)) > 0$. Let now $\Delta < \epsilon'_{kL}/\epsilon'_{kH} \leq \epsilon''_{kL}/\epsilon''_{kH}$ and $\epsilon'_{kL} + \epsilon''_{kL} \leq \epsilon'_{kH} + \epsilon''_{kH}$. As the first two inequalities ensure, respectively, that $\lambda_{L\hat{n}}^k(m_L)/\lambda_{H\hat{n}}^k(m_L) > p(\hat{m})/[1-p(\hat{m})]$ and $\lambda_{H\hat{n}}^k(m_L) = 1 - \lambda_{L\hat{n}}^k(m_L) \geq 1 - \lambda_{L\hat{n}}^k(m_H) = \lambda_{H\hat{n}}^k(m_H)$ everywhere along the sequences, we have

$$\begin{aligned} \Pi_{nk}(\hat{s}_{\hat{n}}, s_{-\hat{n}}^*, \alpha) &= \lambda_{L\hat{n}}^k(m_L) \Pi_L(\mathbf{a}_L(m_L)) + \lambda_{H\hat{n}}^k(m_H) \Pi_H(\mathbf{a}_H(m_H)) \\ &\leq \lambda_{L\hat{n}}^k(m_L) \Pi_L(\mathbf{a}_L(m_L)) + \lambda_{H\hat{n}}^k(m_L) \Pi_H(\mathbf{a}_H(m_H)) \\ &= \lambda_{L\hat{n}}^k(m_L) \left[\Pi_L(\mathbf{a}_L(m_L)) + \frac{\lambda_{H\hat{n}}^k(m_L)}{\lambda_{L\hat{n}}^k(m_L)} \Pi_H(\mathbf{a}_H(m_H)) \right] \\ &< \lambda_{L\hat{n}}^k(m_L) \left[\Pi_L(\mathbf{a}_L(m_L)) + \frac{1-p(\hat{m})}{p(\hat{m})} \Pi_H(\mathbf{a}_H(m_H)) \right] = 0 \end{aligned}$$

The reasoning for why it is optimal for any $n^* \in \mathcal{N}$ to honor m at stage 3 is the same as in Step 1(i) of the proof for Proposition 3.1. ■

Lemma B.2 *Let $(s^*, \alpha^*) \in S_F \times S_A \times S_A$ be an equilibrium profile for Γ when $\mathcal{C} = \{C\}$. Then, $\{\alpha_L(\mathcal{M}(s^*)), \alpha_H(\mathcal{M}(s^*))\}$ is separating and $\Pi_h(\alpha_L(\mathcal{M}(s^*))) = 0$ for either h .*

Proof. Let $\mathbf{a}_h^* = \alpha_h(\mathcal{M}(s^*))$ for either h . Suppose first that $\Pi_L(\mathbf{a}_L^*) = \delta > 0$ and consider the contract $\hat{\mathbf{a}}_L = \mathbf{a}_L^* + (1, \kappa_L)\epsilon_L$ for some $\kappa_L \in (I_H(\mathbf{a}_L^*), I_L(\mathbf{a}_L^*))$ and $\epsilon_L < 0$. This gives

$$\begin{aligned} \Pi_L(\hat{\mathbf{a}}_L) &= (1-p_L)(a_{0L}^* + \epsilon_L) - p_L(a_{1L}^* + \kappa_L\epsilon_L) = \Pi_L(\mathbf{a}_L^*) + [1-p_L(1+\kappa_L)]\epsilon_L \\ &= \delta + [1-p_L(1+\kappa_L)]\epsilon_L \end{aligned}$$

and sufficiently small $|\epsilon_L|$ ensures that $|1-p_L(1+\kappa_L)||\epsilon_L| < \delta$ or $\Pi_L(\hat{\mathbf{a}}_L) > 0$. Yet, by Lemma 1.3 in the Online Appendix, we also have $\hat{\mathbf{a}}_L \succ_L \mathbf{a}_L^* \succ_H \hat{\mathbf{a}}_L$. Which, since $\mathbf{a}_L^* \succ_L \mathbf{a}_H^* \succ_H \mathbf{a}_L^*$, gives $\hat{\mathbf{a}}_L \succ_L \mathbf{a}_H^* \succ_H \hat{\mathbf{a}}_L$. The pooling menu $\{\hat{\mathbf{a}}_L, \hat{\mathbf{a}}_L\}$ attracts only the low-risk type away from $\{\mathbf{a}_L^*, \mathbf{a}_H^*\}$ and makes strictly positive profits doing so. A similarly absurd conclusion follows from the hypothesis $\Pi_H(\mathbf{a}_H^*) > 0$ by deploying the contract $\hat{\mathbf{a}}_H = \mathbf{a}_H^* + (1, \kappa_H)\epsilon_H$ for $\kappa_H \in (I_H(\mathbf{a}_L^*), I_L(\mathbf{a}_L^*))$ and sufficiently small $\epsilon_H > 0$. This shows that it cannot be $\Pi_h(\mathbf{a}_h^*) \leq 0$ for either h . That $\Pi_h(\mathbf{a}_h^*) = 0$ for either h follows then from free entry.

It remains to rule out that $\{\mathbf{a}_L^*, \mathbf{a}_H^*\}$ is pooling. To this end, observe first that $\mathbf{a}_L^* \sim_h \mathbf{a}_H^*$ for either h implies that $\mathbf{a}_L^* = \mathbf{a}_H^*$ (see Lemma 1.8 in the Online Appendix). Similarly, if $\mathbf{a}_L^* \succ_h \mathbf{a}_H^*$ [resp. $\mathbf{a}_H^* \succ_h \mathbf{a}_L^*$] for either h , we can replace $\{\mathbf{a}_L^*, \mathbf{a}_H^*\}$ with $\{\mathbf{a}_L^*, \mathbf{a}_L^*\}$ [resp. with $\{\mathbf{a}_H^*, \mathbf{a}_H^*\}$] without loss of generality. It suffices, therefore, to rule out that $\mathbf{a}_L^* = \mathbf{a}_H^*$. Let then $\mathbf{a}_L^* = \mathbf{a}_H^* = \mathbf{a}^*$ and observe that, since $p_H > p_L$, it cannot be $\Pi_L(\mathbf{a}^*) = 0 = \Pi_H(\mathbf{a}^*)$ unless $\mathbf{a}^* = \mathbf{0}$. As however $\mathbf{a}_h^{**} \succ_h \mathbf{0}$ (see Lemma 2.5(i) in the Online Appendix) while $\mathcal{M}(s^*) = \mathcal{M}(s^{1*})$ when $\mathcal{C} = \{C\}$, a contradiction follows from Lemma A.1. ■

Lemma B.3 *Let $\mathbf{a} \in \mathcal{A}$ be such that $\mathbf{a} \succ_L \mathbf{a}_L^{**}$ while $\Pi(\{\mathbf{a}, \mathbf{a}\}|\lambda) = 0$. Then $\mathbf{a} \in \mathbb{R}_{++}^2$.*

Proof. Since $\mathbf{a} \neq \mathbf{0}$ (see Lemma 2.5(i) in the Online Appendix), $\Pi(\{\mathbf{a}, \mathbf{a}\}|\lambda) = 0$ requires that $\mathbf{a} \in \mathbb{R}_{++}^2 \cup \mathbb{R}_{--}^2$. And as also $\Pi_L(\mathbf{a}_L^{**}) = 0$, letting $\bar{p} = \lambda p_L + (1-\lambda)p_H$ gives $a_{0L}^{**} = (1-p_L)/p_L a_{1L}^{**}$

and $a_0 = (1 - \bar{p})/\bar{p}a_1$. To establish now the claim by contradiction, suppose that $\mathbf{a} \in \mathbb{R}_{-}^2$. Since $\mathbf{a}_L^{**} \in \mathbb{R}_{++}^2$ (see Lemma 2.5(iv) in the Online Appendix), it cannot be $a_{sL}^{**} = a_s + \epsilon_s$ with $\epsilon_s > 0$ for $s \in \{0, 1\}$. Which gives

$$\epsilon_0 = a_{0L}^{**} - a_0 = \frac{1 - p_L}{p_L} a_{1L}^{**} - \frac{1 - \bar{p}}{\bar{p}} a_1 < \frac{1 - p_L}{p_L} (a_{1L}^{**} - a_1) = \frac{1 - p_L}{p_L} \epsilon_1$$

the inequality because $p_L < \bar{p}$. Letting, therefore, $\kappa = \epsilon_0/\epsilon_1$ we have $\mathbf{a}_L^{**} = \mathbf{a} + (\kappa, 1) \epsilon_1$ for some $\kappa \in \left(0, \frac{1 - p_L}{p_L}\right)$. But as \mathbf{a}_L^{**} offers under-insurance, this necessitates (see Lemma 1.2 in the Online Appendix) that $\mathbf{a}_L^{**} \succ_L \mathbf{a}$. ■

Proof of Claim 4.3

To establish the “only if” direction of the first part of the claim, we will argue by contradiction. To this end, suppose (without loss of generality) that $\exists \mathbf{a} \in \mathcal{A}$ s.t. $\mathbf{a} \succ_L \mathbf{a}_L^{**}$ and $\Pi(\{\mathbf{a}, \mathbf{a}\}|\lambda) = 0$.³⁷ Since $\mathbf{a} \succ_L \mathbf{a}_L^{**} \succ_L \mathbf{0}$, the continuity of the preference relation \succ_L requires the existence of some $\pi \in (0, 1)$ such that $\pi \mathbf{a} + (1 - \pi) \mathbf{0} = \pi \mathbf{a} := \hat{\mathbf{a}} \sim_L \mathbf{a}_L^{**}$.³⁸ And as obviously $\Pi(\{\hat{\mathbf{a}}, \hat{\mathbf{a}}\}|\lambda) = 0$ while $\pi < 1$ and $\mathbf{a} \in \mathbb{R}_{++}^2$ (Lemma B.3), letting $\bar{p} = \lambda p_L + (1 - \lambda) p_H$ it must be $\hat{\mathbf{a}} = \mathbf{a} - \left(\frac{\bar{p}}{1 - \bar{p}}, 1\right) \epsilon$ for some $\epsilon > 0$. By Lemma 1.1 in the Online Appendix then we have

$$\begin{aligned} U_L(\mathbf{a}) - U_L(\hat{\mathbf{a}}) &= \left[p_L u'(\hat{w}_1 + \hat{\epsilon}) - \frac{\bar{p}(1 - p_L)}{1 - \bar{p}} u' \left(\hat{w}_0 - \frac{\bar{p}\hat{\epsilon}}{1 - \bar{p}} \right) \right] \epsilon \\ &= \left[p_L u'(\tilde{w}_1) - \frac{\bar{p}(1 - p_L)}{1 - \bar{p}} u'(\tilde{w}_0) \right] \epsilon \end{aligned}$$

for some $\hat{\epsilon} \in (0, \epsilon)$, and where we let $(\hat{w}_0, \hat{w}_1) = w(\hat{\mathbf{a}})$, $\tilde{\mathbf{a}} = \hat{\mathbf{a}} + \left(\frac{\bar{p}}{1 - \bar{p}}, 1\right) \hat{\epsilon}$, and $(\tilde{w}_0, \tilde{w}_1) = w(\tilde{\mathbf{a}})$.

Observe now that $\mathbf{a} \succ_L \mathbf{a}_L^{**} \sim_L \hat{\mathbf{a}}$ implies $\frac{\bar{p}}{1 - \bar{p}} < \frac{p_L u'(\tilde{w}_1)}{(1 - p_L) u'(\tilde{w}_0)} = I_L(\tilde{\mathbf{a}})^{-1}$. We may consider thus the contract $\mathbf{a}' = \tilde{\mathbf{a}} + (\kappa, 1) \epsilon$ with $\kappa \in \left(\frac{\bar{p}}{1 - \bar{p}}, I_L(\tilde{\mathbf{a}})^{-1}\right)$ and $\epsilon > 0$. Since $\Pi(\{\tilde{\mathbf{a}}, \tilde{\mathbf{a}}\}|\lambda) = 0$, as a pooling menu this new contract gives

$$\Pi(\{\mathbf{a}', \mathbf{a}'\}|\lambda) = \Pi(\{\tilde{\mathbf{a}}, \tilde{\mathbf{a}}\}|\lambda) + [\kappa(1 - \bar{p}) - \bar{p}] \epsilon = [\kappa(1 - \bar{p}) - \bar{p}] \epsilon > 0$$

For either h , moreover, applying again Lemma 1.1 in the Online Appendix successively we get

$$\begin{aligned} U_h(\mathbf{a}') - U_h(\hat{\mathbf{a}}) &= U_h(\mathbf{a}') - U_h(\tilde{\mathbf{a}}) + U_h(\tilde{\mathbf{a}}) - U_h(\hat{\mathbf{a}}) \\ &= [p_h u'(\tilde{w}_1 + \tilde{\epsilon}) - \kappa(1 - p_h) u'(\tilde{w}_0 - \kappa \tilde{\epsilon})] \epsilon \\ &\quad + \left[p_h u'(\hat{w}_1 + \epsilon') - \frac{\bar{p}(1 - p_h)}{1 - \bar{p}} u' \left(\hat{w}_0 - \frac{\bar{p}\epsilon'}{1 - \bar{p}} \right) \right] \hat{\epsilon} \\ &= [\kappa^{-1} - I_h(\tilde{\mathbf{a}} + (\kappa, 1) \tilde{\epsilon})] p_h u'(\tilde{w}_1 + \tilde{\epsilon}) \kappa \epsilon \\ &\quad + \left[p_h u'(\hat{w}_1 + \epsilon') - \frac{\bar{p}(1 - p_h)}{1 - \bar{p}} u' \left(\hat{w}_0 - \frac{\bar{p}\epsilon'}{1 - \bar{p}} \right) \right] \hat{\epsilon} \end{aligned}$$

³⁷If $\mathbf{a} \in \mathcal{A}$ is such that $\Pi(\{\mathbf{a}, \mathbf{a}\}|\lambda) = \epsilon > 0$ then the contract $\hat{\mathbf{a}} = \mathbf{a} - (1, -1) \epsilon$ gives $\Pi(\{\hat{\mathbf{a}}, \hat{\mathbf{a}}\}|\lambda) = \Pi(\{\mathbf{a}, \mathbf{a}\}|\lambda) - \epsilon = 0$ and $\hat{\mathbf{a}} \succ_h \mathbf{a}$ for either h .

³⁸Recall that there is a one-to-one relation between contract and income points. The continuity of the relation \succ_L derives from continuous preferences over lotteries on wealth vectors. Given this, the existence of the wealth vector in question (and, thus, of the corresponding contract) is a standard result to be found in textbook derivations of the expected utility theorem (see for instance Step 3 of Proposition 6.B.3 in Mas-Colell et al. [24]).

for some $(\tilde{\varepsilon}, \epsilon') \in (0, \varepsilon) \times (0, \hat{\varepsilon})$. Yet, $\kappa^{-1} > I_L(\tilde{\mathbf{a}}) > I_H(\tilde{\mathbf{a}})$ and, letting $\Delta = \kappa^{-1} - I_L(\tilde{\mathbf{a}})$, we may choose ε (and subsequently $\tilde{\varepsilon}$) small enough to guarantee that $|I_h(\tilde{\mathbf{a}} + (\kappa, 1)\tilde{\varepsilon}) - I_h(\tilde{\mathbf{a}})| < \Delta$ for either h . But then for either h

$$\begin{aligned}
U_h(\mathbf{a}') - U_h(\hat{\mathbf{a}}) &= [\Delta + I_L(\tilde{\mathbf{a}}) - I_h(\tilde{\mathbf{a}} + (\kappa, 1)\tilde{\varepsilon})] p_h u'(\tilde{w}_1 + \tilde{\varepsilon}) \kappa \varepsilon \\
&\quad + \left[p_h u'(\hat{w}_1 + \epsilon') - \frac{\bar{p}(1-p_h)}{1-\bar{p}} u' \left(\hat{w}_0 - \frac{\bar{p}\epsilon'}{1-\bar{p}} \right) \right] \hat{\varepsilon} \\
&> \left[p_h u'(\hat{w}_1 + \epsilon') - \frac{\bar{p}(1-p_h)}{1-\bar{p}} u' \left(\hat{w}_0 - \frac{\bar{p}\epsilon'}{1-\bar{p}} \right) \right] \hat{\varepsilon} \\
&> \left[p_h u'(\hat{w}_1 + \hat{\varepsilon}) - \frac{\bar{p}(1-p_h)}{1-\bar{p}} u' \left(\hat{w}_0 - \frac{\bar{p}\hat{\varepsilon}}{1-\bar{p}} \right) \right] \hat{\varepsilon} \\
&\geq \left[p_L u'(\hat{w}_1 + \hat{\varepsilon}) - \frac{\bar{p}(1-p_L)}{1-\bar{p}} u' \left(\hat{w}_0 - \frac{\bar{p}\hat{\varepsilon}}{1-\bar{p}} \right) \right] \hat{\varepsilon} \\
&= [U_L(\mathbf{a}) - U_L(\hat{\mathbf{a}})] \frac{\hat{\varepsilon}}{\varepsilon} > 0
\end{aligned}$$

where the second inequality follows from the fact that $\epsilon' < \hat{\varepsilon}$ while $u(\cdot)$ is strictly concave; by contrast, the last inequality exploits that $p_H > p_L$.

Let now the menu $\{\mathbf{a}'\}$ be offered in the presence of the RSW. As $\mathbf{a}' \succ_L \hat{\mathbf{a}} \sim_L \mathbf{a}_L^{**} \succ_L \mathbf{a}_H^{**}$, the low-risk type is pulled away. If $\mathbf{a}' \succ_H \mathbf{a}_H^{**}$ so is the high-risk type, rendering the deviant menu strictly-profitable as pooling. Otherwise, the deviant attracts only the low-risk type delivering even higher expected profits - since $\mathbf{a} \in \mathbb{R}_{++}^2$, choosing ε sufficiently small guarantees that $\hat{\mathbf{a}}, \tilde{\mathbf{a}}, \mathbf{a}' \in \mathbb{R}_{++}^2$; that $\Pi_L(\mathbf{a}') > \Pi(\{\mathbf{a}', \mathbf{a}'\}|\lambda)$ follows then from $p_L < \bar{p}$.

For the “if” direction of the first part of the claim, recall first the argument that established Claim 4.2 in the main text. Credible challenges to the RSW allocation may come only from menus $\hat{m} \in \mathcal{M}$ such that $\mathbf{a}_h(\hat{m}) \succ_h \mathbf{a}_h^{**}$ for either h while $\Pi(\hat{m}|\lambda) > 0$. Given though that now we restrict attention to the space of safe separating menus, in which the RSW allocation is dominant for either risk-type, credible challenges can come only from pooling menus $\{\hat{\mathbf{a}}, \hat{\mathbf{a}}\} \in \mathcal{M}$: $\hat{\mathbf{a}} \succ_h \mathbf{a}_h^{**}$ for either h while $\Pi(\{\hat{\mathbf{a}}, \hat{\mathbf{a}}\}|\lambda) > 0$. By hypothesis however there are no such menus (at least for the low-risk type).

For the second part of the claim, take an arbitrary $(u, p_L, p_H) \in \mathcal{U} \times (0, 1)^2$ and consider the function $\tilde{\mathbf{a}} : (0, 1) \rightarrow \mathbb{R}_{++}^2$ given by $\tilde{\mathbf{a}}(\lambda) = \arg \max_{\mathbf{a} \in \mathcal{A}: \Pi(\{\mathbf{a}, \mathbf{a}\}|\lambda) \geq 0} U_L(\mathbf{a})$. As the underlying optimization problem entails an objective function that is strictly-concave and continuous on a constraint set that for each $\lambda \in (0, 1)$ is compact, $\tilde{\mathbf{a}}(\cdot)$ is a well-defined solution function on $(0, 1)$. As moreover the IE(λ)-optimum $\{\mathbf{a}^\lambda\}$ satisfies the constraint while $\mathbf{a}^\lambda \succ_L \mathbf{a}_L^{**}$, it must be $\tilde{\mathbf{a}}(\lambda) \succ_L \mathbf{a}_L^{**}$ for any $\lambda \in (0, 1)$. That $\tilde{\mathbf{a}}(\lambda) \in \mathbb{R}_{++}^2$ for any $\lambda \in (0, 1)$ follows then from the fact that the constraint must bind at the optimum (recall footnote 37) and Lemma B.3. The balance of our argument will consist of the following observations.

(i) The value function $U_L(\tilde{\mathbf{a}}(\cdot))$ is strictly increasing on $(0, 1)$.

To see this, notice that $p_H > p_L$, $\tilde{\mathbf{a}}(\lambda) \in \mathbb{R}_{++}^2$, and $\Pi(\{\tilde{\mathbf{a}}\}(\lambda)|\lambda) = 0$ together imply that $\Pi_L(\tilde{\mathbf{a}}(\lambda)) > 0 > \Pi_H(\tilde{\mathbf{a}}(\lambda))$ for any $\lambda \in (0, 1)$. Letting then $\lambda, \lambda' \in (0, 1) : \lambda' < \lambda$ we must

have

$$\begin{aligned}
\lambda \Pi_L(\tilde{\mathbf{a}}(\lambda)) + (1 - \lambda) \Pi_H(\tilde{\mathbf{a}}(\lambda)) &= 0 \\
&= \lambda' \Pi_L(\tilde{\mathbf{a}}(\lambda')) + (1 - \lambda') \Pi_H(\tilde{\mathbf{a}}(\lambda')) \\
&< \lambda \Pi_L(\tilde{\mathbf{a}}(\lambda')) + (1 - \lambda) \Pi_H(\tilde{\mathbf{a}}(\lambda'))
\end{aligned}$$

That is $\tilde{\mathbf{a}}(\lambda) \neq \tilde{\mathbf{a}}(\lambda')$, and thus $U_L(\tilde{\mathbf{a}}(\lambda')) < U_L(\tilde{\mathbf{a}}(\lambda))$ given that $\tilde{\mathbf{a}}(\cdot)$ is a function.

(ii) $U_L(\tilde{\mathbf{a}}(\cdot))$ being strictly monotone on $(0, 1)$, the parameter

$$\lambda_0 \in (0, 1) : U_L(\tilde{\mathbf{a}}(\lambda_0)) = U_L(\mathbf{a}_L^{**})$$

is defined uniquely. Consider then the canonical insurance economy with parameters (u, p_L, p_H, λ_0) . Condition (8) being satisfied, the RSW menu is the equilibrium in this economy. However, the pooling menu $\{\tilde{\mathbf{a}}(\lambda_0)\}$ being feasible but obviously not optimal in the MWS-problem, it cannot but be $\bar{\mathbf{a}}_L(\lambda_0) \succ_L \tilde{\mathbf{a}}(\lambda_0) \sim_L \mathbf{a}_L^{**}$ and the RSW allocation is not IE in the economy with parameters (u, p_L, p_H, λ_0) (see Claim 3.1 in the Online Appendix).

(iii) Consider next the function $\bar{\mathbf{a}}_L : (0, 1) \rightarrow \mathbb{R}_{++}^2$ that gives the low-risk MWS-contract in the canonical insurance economy with parameters (u, p_L, p_H, λ) for $\lambda \in (0, 1)$. This function is well-defined while $\Pi_L(\bar{\mathbf{a}}_L(\lambda)) \geq 0 \geq \Pi_H(\bar{\mathbf{a}}_L(\lambda))$ for any for $\lambda \in (0, 1)$ (see Claim 2.1 and Lemmas 2.5(v), 2.4 in the Online Appendix). The respective argument in step (i) above can be deployed again here to establish that $U_L(\bar{\mathbf{a}}_L(\cdot))$ is weakly increasing on $(0, 1)$.

(iv) Let now $\Delta = U_L(\bar{\mathbf{a}}_L(\lambda_0)) - U_L(\tilde{\mathbf{a}}(\lambda_0))$, and consider an arbitrary $\lambda \in (0, \lambda_0)$. Since $U_L(\tilde{\mathbf{a}}(\lambda)) < U_L(\tilde{\mathbf{a}}(\lambda_0)) = U_L(\mathbf{a}_L^{**})$, the RSW allocation remains the equilibrium of the economy with parameters (u, p_L, p_H, λ) for any $\lambda \in (0, \lambda_0)$.

(v) By step (iii) however any $\lambda \in (0, \lambda_0)$ gives $U_L(\bar{\mathbf{a}}_L(\lambda)) \leq U_L(\bar{\mathbf{a}}_L(\lambda_0))$ and, by the continuity of $U_L(\bar{\mathbf{a}}_L(\cdot))$ (see Lemma 2.2 in the Online Appendix), there exists $\lambda_1 \in [0, \lambda_0)$ s.t. $U_L(\bar{\mathbf{a}}_L(\lambda_0)) - U_L(\bar{\mathbf{a}}_L(\lambda)) < \Delta$ for all $\lambda \in (\lambda_1, \lambda_0)$. It follows then that

$$\begin{aligned}
U_L(\bar{\mathbf{a}}_L(\lambda)) - U_L(\mathbf{a}_L^{**}) &= U_L(\bar{\mathbf{a}}_L(\lambda)) - U_L(\bar{\mathbf{a}}_L(\lambda_0)) \\
&+ U_L(\bar{\mathbf{a}}_L(\lambda_0)) - U_L(\tilde{\mathbf{a}}(\lambda_0)) + U_L(\tilde{\mathbf{a}}(\lambda_0)) - U_L(\mathbf{a}_L^{**}) \\
&= U_L(\bar{\mathbf{a}}_L(\lambda)) - U_L(\bar{\mathbf{a}}_L(\lambda_0)) + \Delta > 0 \quad \forall \lambda \in (\lambda_1, \lambda_0)
\end{aligned}$$

The RSW allocation is not IE in the economy (u, p_L, p_H, λ) for any $\lambda \in (\lambda_1, \lambda_0]$. ■