Insurance Mathematics and Economics 116 (2024) 74–94

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

journal homepage: www.elsevier.com/locate/ime

Can price collars increase insurance loss coverage?

Indradeb Chatterjee, MingJie Hao, Pradip Tapadar *, R. Guy Thomas

School of Mathematics, Statistics and Actuarial Science, University of Kent, Canterbury, CT2 7NF, UK

A R T I C L E I N F O

JEL classification:
D82
G22

Keywords:
Insurance loss coverage
Risk classification
Price collar
Demand elasticity
Marginal revenue-to-cost ratio
Arc elasticity
Partial community rating

A B S T R A C T

Loss coverage, defined as expected population losses compensated by insurance, is a public policy criterion for comparing different risk-classification regimes. Using a model with two risk-groups (high and low) and iso-elastic demand, we compare loss coverage under three alternative regulatory regimes: (i) full risk-classification (ii) pooling (iii) a price collar, whereby each insurer is permitted to set any premiums, subject to a maximum ratio of its highest and lowest prices for different risks. Outcomes depend on the comparative demand elasticities of low and high risks. If low-risk elasticity is sufficiently low compared with high-risk elasticity, pooling is optimal; and if it is sufficiently high, full risk-classification is optimal. For an intermediate region where the elasticities are not too far apart, a price collar is optimal, but only if both elasticities are greater than one. We give extensions of these results for more than two risk-groups. We also outline how they can be applied to other demand functions using the construct of arc elasticity.

1. Introduction

Restrictions on risk classification are common in retail insurance markets, usually in the form of bans on the use of particular variables (e.g. gender, race, genetic tests). Restrictions of this type can align insurance practice with social norms which deprecate discrimination on particular variables. But the use of deprecated variables is not the only public policy issue with risk classification. Another concern is the equity of a wide dispersion in prices for different risks, irrespective of the particular underlying variables which it might reflect; and in particular, the problem of unaffordable insurance for higher risks. For this type of concern, piecemeal restriction of specific variables seems an indirect and hard-to-calibrate form of regulation. Also, if insurers can use other variables which are correlated with the banned ones, the effect of the regulation can be partly undone, unless more prescriptive regulation requires estimation procedures which neutralise this effect (Pope and Sydnor (2011), Lindholm et al. (2021)).

A potentially more direct way of addressing equity concerns is to set an explicit limit on the dispersion of prices for different risks, via what we call a *price collar*. A price collar (sometimes called "partial community rating") is a maximum ratio, say κ, for the highest and lowest prices an insurer may charge for different risks. A price collar restricts only an insurer's relative prices, not the absolute level. It can be used either in addition to, or instead of, limits on specific variables. One example is the Affordable Care Act in the US, which permits differentiation by a factor of up to 1.5× for tobacco use and 3× for age (and no other factors except coverage tier, geography, and number of dependants).

In this paper, we abstract from the reference to specific variables and consider a price collar which limits the overall dispersion of an insurer's prices. The collar regime, which can also be described as *partial risk-classification*, represents a compromise approach, in between *full risk-classification* (where prices are fully differentiated for individual risks, which we assume are fully observable via insurers' underwriting procedures), and *pooling* (where all risk classification is banned and each insurer sets a single price, the same for all risks).

1.1. Loss coverage

A price collar can potentially ameliorate equity concerns, but policymakers also need to consider its efficiency consequences, in particular the adverse selection which may be induced. We take a nuanced view of this: we argue that whilst excessive adverse selection is a concern, a modest degree of adverse selection can actually increase efficiency, if it increases what we call "loss coverage". This concept has been described elsewhere

* Corresponding author.
E-mail address: P.Tapadar@kent.ac.uk (P. Tapadar).

https://doi.org/10.1016/j.insmatheco.2024.02.003
Received February 2023; Received in revised form October 2023; Accepted 7 February 2024
Available online 15 February 2024

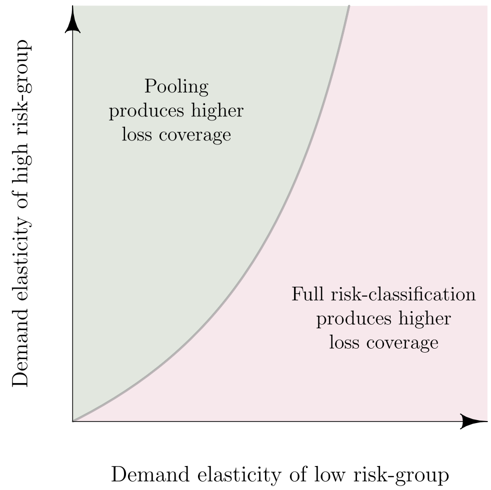

Fig. 1. Regions where loss coverage is higher for pooling and full risk-classification.

(Thomas (2008, 2009, 2017); Hao et al. (2016, 2018, 2019), but it may remain unfamiliar to many readers; so we now review the rationale, in the specific context of a price collar.

The usual story about adverse selection runs as follows. Compared to fully risk-differentiated prices, the imposition of a price collar raises prices for low risks and lowers prices for high risks. Low risks buy less insurance, and high risks buy more (adverse selection). The weighted average price paid for insurance across the whole market rises. Also, because high risks are typically less numerous than low risks, the number of risks insured falls. This combination of a rise in average price and fall in demand is usually seen as a bad outcome, both for insurers and for society.

However, the social purpose of insurance is to compensate the population's losses (not to sell "coverage" indiscriminately, with no regard for any prospect of loss). Insurance of one high risk contributes more in expectation to this objective than insurance of one low risk. This suggests that public policymakers might welcome increased purchasing by high risks, except for the usual story about adverse selection.

The usual story about adverse selection overlooks one point: with adverse selection, expected losses compensated by insurance – a quantity we term "loss coverage" – can be *higher* than with no adverse selection. The rise in weighted average price when a price collar is imposed reflects a shift in coverage towards higher risks. From a public policymaker's viewpoint, this means that more of the "right" risks (i.e. those more likely to suffer loss) buy insurance. If this shift in coverage is large enough, it can more than offset the fall in numbers insured, so that loss coverage is increased. We argue that where adverse selection leads to higher loss coverage – that is, more risk being voluntarily transferred and more losses being compensated – then from society's perspective, this should be seen as a good outcome from adverse selection. Formally, we define loss coverage as:

$$\text{Loss coverage} = \text{Expected losses compensated by insurance} \tag{1.1}$$

and then we argue that a policymaker should prefer risk-classification schemes which generate higher loss coverage for the population as a whole (note: not necessarily the same as higher number of persons insured, irrespective of their individual risks).

Another way of putting this is that a public policymaker designing a risk-classification scheme in the context of adverse selection faces a trade-off between insurance of the "right" risks (those more likely to suffer loss), and insurance of a larger number of risks. The optimal trade-off depends on demand elasticities of high and low risks. The concept of loss coverage quantifies this trade-off, and provides a metric for comparing the efficacy of different risk-classification schemes in facilitating compensation for the losses of the population as a whole.[1]

Previous work (Hao et al. (2016, 2018)) used models of demand elasticity for high and low risks to compare loss coverage for the polar cases: full risk-classification (where prices fully reflect individual risks) and pooling (where all risk classification is banned). This led to results as sketched in Fig. 1. For some intuition on this pattern, suppose we start from full risk-classification and then increase the price for low risks and reduce the price for high risks (i.e. move towards pooling), while maintaining zero profits. If demand elasticity for low risks is low compared to that for high risks, only a small fraction of low risks leaves the market, while a large fraction of high risks enters the market. Furthermore, each high risk who enters "counts for more" towards our policy objective of expected losses compensated than each low risk who leaves. So even though coverage of *persons* typically falls on a move towards pooling,[2] coverage of *losses* can rise.

Given the pattern of the two regions in Fig. 1, it is natural to wonder if there might be an intermediate region where a price collar – a compromise between pooling and full risk-classification – gives higher loss coverage than either of the two polar cases. That possibility is the subject of this paper. To anticipate our main conclusion, the answer is yes, but only if demand elasticities for both high and low risks are greater than 1.

1.2. Literature review

This paper is related to recent literature on State-level health insurance exchanges established in the United States under the Affordable Care Act (and earlier legislation in some States), which all involve some form of price collar (Ericson and Starc (2015), Mahoney and Weyl (2017), Einav et al. (2019), Geruso et al. (2023)). These papers all focus on the specific context of US healthcare, and so incorporate many institutional constraints and details which we omit: a fixed width of collar (in contrast, our primary focus is on the effect of varying the collar), risk adjustment,

[1] For a toy example illustrating the arithmetic of loss coverage for a small group of risks, see section 2 of Hao et al. (2018) or this link: Improving insurance with some adverse selection.

[2] Coverage of *persons* typically falls on a move towards pooling, despite the posited difference in elasticities, because the high risk-group represents a smaller fraction of the population than the low risk-group. But we argue that this is not a disadvantage, if coverage of *losses* rises.

premium subsidies for insurance purchase and penalties for non-purchase, etc. One of the conclusions is that demand elasticity on the healthcare exchanges for younger people (i.e. lower risks) appears empirically to be about twice that for older people (i.e. higher risks). Therefore if subsidies were differentiated by age to make premiums lower for younger people, this could increase enrolment (Ericson and Starc (2015)), and represents a more effective and flexible instrument for this purpose than risk adjustment (Einav et al. (2019)). This is broadly consistent with our results: if demand elasticity is high for low risks compared with high risks, more differentiated prices tend to increase coverage (and loss coverage). However the details of our analysis differ, in that we assume the policy criterion is to maximise loss coverage (risk-weighted demand), not numbers insured (un-weighted demand).

To illustrate the difference, if each high risk has two times the expected losses of each low risk, then on our criterion, the policymaker is indifferent between coverage of one high risk or two low risks; but in the other papers just cited, the policymaker is indifferent between coverage of one high risk and one low risk. In the healthcare example, our objective amounts to increasing (expected) coverage of sickness, rather than coverage of persons irrespective of their probability of sickness. We do not say that prioritising coverage of sickness over coverage of persons is the only reasonable preference, but we do say that it is at least arguable.

Einav et al. (2019) recognise the point when they note that their suggestion of a move from risk adjustment and uniform premium subsidies to less risk adjustment and more risk-differentiated subsidies, whilst increasing enrolment, could also make higher risks worse off; to address this, they suggest adding a Pareto-type restriction on the policy change (i.e. the effective prices faced by high risks, after the move to less risk adjustment and more differentiated subsidies, must be no higher than before the change). The restriction ensures that the policy change increases enrolment for low risks, and also does not reduce enrolment for high risks; in our terms, it ensures that the policy change increases loss coverage. But it does not necessarily *maximise* loss coverage.

The contribution of this paper is to analyse market outcomes over the full parameter space for demand elasticities of high and low risk-groups (iso-elastic for each risk-group); and over the full feasible range for a price collar, from complete pooling to full risk-classification. The paper also represents a natural extension of Hao et al. (2018), which considered only the polar cases of full risk-classification and pooling, and not the intermediate possibilities under a price collar.

2. Political and regulatory constraints

By “political constraints”, we mean certain general notions of fairness and proportionality, which are unlikely to be explicitly stated in insurance regulations, but may nevertheless constrain both the actions of regulators and insurers’ response to those actions. By “regulatory constraints”, we mean rules about risk-classification which are explicitly stated in insurance regulations (such as the rules imposing a price collar).

2.1. Political constraints

The main role of “political constraints” in our analysis is to rule out premium regimes which may be technically feasible, but which we think are likely to be widely regarded as perverse, unreasonable, or otherwise unacceptable. Here are some examples:

(i) *Hyper-differentiation.* High risks are charged more than their true risks, and low risks are charged less than theirs. This gives a cross-subsidy between risk-groups, but for most contexts, it is in the “wrong” direction (i.e. a disadvantaged high risk-group is over-charged to subsidise an already fortunate low risk-group).
(ii) *Hyper-pooling.* Low risks are charged more than high risks. This could be economically feasible, if demand elasticity for low risks is sufficiently low compared to that for high risks. Here, the cross-subsidy is in the “right” direction for equity, but too large.
(iii) *Groupings which lead to unfair ordering of premiums.* If there are three risk-groups, say low, medium and high risks, then a regime which groups low and high risk-groups at one premium, and medium risks at another premium, could be economically feasible. But this might be politically unacceptable, if it leads to low risks being charged more than medium risks. (Stated differently, the objection is that this gives a similar result to hyper-pooling of low and medium risks.)

To summarise these considerations: if premiums are differentiated, the span of premiums should not exceed the span of the risks, and the ordering of premiums should be the same as the ordering of risks.[3] To state these notions formally, we now introduce some notation for the rest of the paper:

- We assume a population consisting of n distinct risk-groups with probabilities of loss given by $\underline{\mu} = (\mu_1, \mu_2, \dots, \mu_n)$, which are ordered for convenience (μ_1 smallest, μ_n largest).
- The proportion of the population belonging to risk-group i is p_i, for $i = 1, 2, \dots, n$.
- Members of risk-group i are offered premium (per unit of loss) π_i. We call $\underline{\pi} = (\pi_1, \pi_2, \dots, \pi_n)$ a premium regime, or risk-classification regime.

The political constraints can then be encapsulated as:

Constraint 1 *(Political). Given risks $\underline{\mu}$, a politically acceptable premium regime $\underline{\pi}$ needs to satisfy:*

$$\mu_1 \le \pi_1 \le \pi_2 \le \cdots \le \pi_n \le \mu_n. \tag{2.1}$$

Other examples of politically unacceptable premium regimes might include those which lack face validity (e.g. combine risk-groups having no apparent similarities), or which disadvantage socially protected classes (e.g. a high premium for high risk may be less acceptable for disability than

[3] The latter constraint is analogous to the principle of vertical equity in taxation: the ordering of post-tax incomes should be the same as the ordering of pre-tax incomes.

for dangerous sports). It is not possible to specify all the political considerations that might arise. But we think Constraint 1 above seems broadly applicable, and helps to rule out perverse interpretations of the price collar concept.

2.2. Regulatory constraints

The price collar concept was outlined in the introduction and is formally specified as Constraint 2.

Constraint 2 *(Price collar). Given a prescribed price collar, κ, where $\kappa \geq 1$, any premium regime $\underline{\pi}$ needs to satisfy:*

$$\pi_H \leq \kappa\, \pi_L, \tag{2.2}$$

where $\pi_L = \min_i \pi_i$ and $\pi_H = \max_i \pi_i$.

The acceptable range for a price collar is limited by the prohibition on hyper-separation and hyper-pooling in political Constraint 1, which implies that:

$$1 \leq \kappa \leq \frac{\mu_n}{\mu_1}. \tag{2.3}$$

Note that the extremes $\kappa = 1$ and $\kappa = \mu_n/\mu_1$ correspond to pooling and full risk-classification respectively. We use the term *partial risk-classification* to refer to all regimes where κ is set to an intermediate value.

A price collar can be circumvented if insurers can simply decline high risks (which amounts to quoting an infinite price above the collar for high risks). To be effective, it needs to be supported by *guaranteed issue*, that is an obligation on the insurer to accept any applicant at some price within the collar, as stated in Constraint 3.

Constraint 3 *(Guaranteed issue). Insurers are required to quote a price within the collar to all applicants. Nobody can be declined for insurance.*

Guaranteed issue might be unreasonable for types of insurance where there are a few exceptionally high probabilities of loss (e.g. term insurance for people with a terminal illness). This problem can be alleviated by a rule which permits a small fraction of higher prices, e.g. up to 1% of the prices charged by an insurer over a trailing three-year period are permitted to exceed the collar. But for simplicity, we assume guaranteed issue is required.

3. Insurance market and loss coverage

This section develops the theory of insurance market equilibrium and loss coverage under perfect competition, but subject to the political and regulatory constraints stated above.[4]

3.1. Insurance demand

Typical theories of insurance demand assume individuals know their own probabilities of loss and have a common utility function. Given an offered premium, individuals with the same probabilities of loss then all make the same purchasing decision. This does not correspond well to the observable reality of insurance markets, where individuals with similar probabilities of loss often appear to make different decisions, and substantial fractions of individuals do not purchase insurance at all.

We follow the different approach introduced in Hao et al. (2018, 2019); Chatterjee et al. (2021), which allows for heterogeneity in risk aversion across individuals with the same probabilities of loss, and hence generates the partial take-up of insurance that we observe in practice. In summary, our approach is based on the following assumptions.

(i) Individuals know their probability of loss and their own risk aversion, and make purchasing decisions accordingly.
(ii) Insurers can observe (via the usual underwriting procedures) individuals' probabilities of loss and so correctly assign them to risk-groups, but cannot observe their individual risk aversion.
(iii) All insurance is for one unit of cover, in a contract which is standardised across all insurers, who compete only on price. Insurers do not offer partial cover or other contract menus.
(iv) Viewed by the insurer, the demand for insurance from risk-group i at premium π_i is then a function $d_i(\pi_i)$. The demand function represents the proportion of the risk-group who buy insurance, such that $0 < d_i(\pi_i) < 1$. We assume that this is decreasing and continuous.

Note that we use demand functions in a purely predictive way, to model how the fraction of a risk-group which purchases insurance varies with the price charged to the risk-group. The role of utility functions is limited to providing the micro-foundation or "back-story" for this predictive model. We do not use the implicit link between utility functions and demand to connect observed changes in demand to changes in social welfare, as in the approach originated by Einav et al. (2010).

[4] Mahoney and Weyl (2017) give an analysis of imperfect competition in selection markets by indexing the degree of competition (and hence mark-ups) with a single parameter. But this cannot be applied in our set-up, where elasticities differ across risk-groups, and hence optimal mark-ups under pooling would vary continuously depending on the exact mix of risks attracted into the pooling equilibrium.

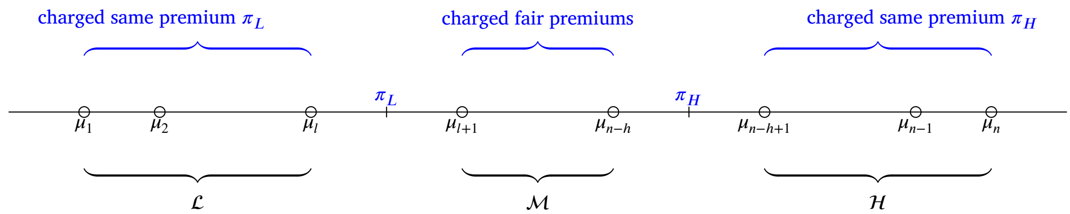

Fig. 2. Price collar: tripartite solution.

3.2. Market equilibrium

To recap on notation: we have a population of n risk-groups, where members of risk-group i each have risk μ_i, are offered premium π_i, and collectively represent proportion p_i of the aggregate population. In a perfectly competitive insurance market, we then have:

$$\text{Premium income} = \sum_{i=1}^{n} p_i\, d_i(\pi_i)\, \pi_i. \tag{3.1}$$

$$\text{(Expected) insurance claim} = \sum_{i=1}^{n} p_i\, d_i(\pi_i)\, \mu_i. \tag{3.2}$$

$$\text{(Expected) profit}: E\left(\underline{\pi}\right) = \sum_{i=1}^{n} p_i\, d_i(\pi_i)\left(\pi_i - \mu_i\right). \tag{3.3}$$

$$\text{Market equilibrium} \Rightarrow E\left(\underline{\pi}\right) = 0. \tag{3.4}$$

Any candidate equilibrium premium regime $\underline{\pi}$ satisfying Equation (3.4) also needs to be a Nash equilibrium: that is, it must be impossible for any single insurer to profitably disrupt the equilibrium by using a different regime, given the way other insurers would react.[5] This property relies on a corollary of perfect competition: *no insurer has any market power*, because all individuals choose the lowest offer in the market for their risk.

To see why this corollary is important, imagine that one insurer could set a collar of the same maximum permitted width κ as other insurers, but with a lower mid-point, and increase its market share of (profit-making) low risks at the lower end of its collar *without* an equivalent increase in market share of (loss-making) high risks at the upper end of its collar. This insurer would profitably disrupt an existing zero-profit regime. But this is impossible under perfect competition, because an insurer which sets lower prices than other insurers for both high and low risks will attract essentially *all* the market demand from both low risks and high risks; and therefore if the other insurers were previously zero-profit, the deviating insurer will make a loss.[6]

3.3. The price collar equilibrium

The equilibrium conditions stated in Section 3.2, combined with the political and regulatory constraints stated in Section 2, lead to the following solutions for different values of the price collar:

(i) $\kappa = \mu_n/\mu_1$: full risk-classification, where all risk-groups are charged their fair premiums i.e. $\underline{\pi} = \underline{\mu}$, is a solution.
(ii) $\kappa = 1$: pooling, where all risk-groups are charged the same premium π_0, is a solution.[7]
(iii) Intermediate values of κ: a tripartite solution of this form:

- a super-group $\mathcal{L}$ of low risk-groups all charged the same π_L (more than their fair premiums);
- a super-group $\mathcal{M}$ of "middle" risk-groups all charged their fair premiums;
- a super-group $\mathcal{H}$ of high risk-groups all charged the same π_H (less than their fair premiums);

where the grouping into three super-groups (hence "tripartite") has the pattern shown in Fig. 2.

For intuition on why the intermediate collar solution must take this tripartite form, first note that each insurer has to come up with its own collar limits, π_L and $\pi_H = \kappa\,\pi_L$. Then all risks lower than π_L have to be charged at least π_L; and all risks higher than π_H have to be charged at most π_H. Competition in pricing for low and high risks drives all insurers towards the same values for the collar limits π_L and π_H. Then note that inside the limits of the collar, the same competitive forces operate as if no collar applied, with the same outcome: if one insurer attempts to pool a higher and lower risk-group inside the collar, this can be destabilised by another insurer offering the lower risk-group in the putative pooling a fair premium. Hence competition drives all insurers to charge all "middle" risk-groups exactly their fair premiums.

[5] Our concept of *different regime* is limited to *different set of prices* for insurance contracts, which are standardised for regulatory or institutional reasons, as in Akerlof (1970); we do not consider differentiation by contract design, as in Rothschild and Stiglitz (1976).

[6] Stating "no market power" differently: *brand elasticity* of demand for each insurer is infinite. Note that no contradiction arises from infinite *brand* demand elasticities combined with relatively low *product* demand elasticities, because they relate to different choice problems. Product elasticity relates to the *choice over products* in the consumer's consumption bundle. Brand elasticity relates to the *choice over brands*, given that insurance is to be purchased. Alternatively, a weaker assumption could do the same work: if one insurer reduces prices for high and low risks by the same amount, it attracts the *same increase in market share* for high and low risks (as assumed in Mahoney and Weyl (2017)); that is, the *semi-logarithmic brand elasticity of demand* is uniform for high and low risks.

[7] Given continuous demand functions, the existence of at least one pooling equilibrium is ensured by the intermediate value theorem. Uniqueness is technically not guaranteed, but pertains for plausible combinations of demand elasticities; and if, exceptionally, there are multiple solutions, any besides the lowest will be eliminated by competition. The same elimination principle applies for any multiple solutions under the collar in the next point (for proof see Appendix A).

A formal statement and proof of the tripartite solution is given as Theorem 1 in Appendix A.

3.4. Iso-elastic insurance demand

So far, we have only needed the insurance demand function, $d_i(\pi_i)$ for risk-group i, to be continuous and decreasing in π_i. We will now further assume differentiability, to define the (point price) elasticity of insurance demand as:

$$\epsilon_i(\pi_i) = -\frac{\partial \log d_i(\pi_i)}{\partial \log \pi_i} = -\frac{\frac{\partial d_i(\pi_i)}{d_i(\pi_i)}}{\frac{\partial \pi_i}{\pi_i}} = -\frac{\pi_i}{d_i(\pi_i)}\frac{\partial d_i(\pi_i)}{\partial \pi_i}, \tag{3.5}$$

which implies that insurance demand can also be expressed as:

$$d_i(\pi_i) = \tau_i \exp\left[-\int_{\mu_i}^{\pi_i} \epsilon_i(s)\, d\log s\right], \tag{3.6}$$

where $\tau_i = d_i(\mu_i)$ is the *fair-premium demand* for risk-group i. A tractable functional form for demand is iso-elastic, i.e. demand elasticity is a positive constant, λ_i:

$$\epsilon_i(\pi_i) = \lambda_i; \tag{3.7}$$

and by Equation (3.6), demand for risk-group i then takes the form:

$$d_i(\pi_i) = \tau_i \left(\frac{\mu_i}{\pi_i}\right)^{\lambda_i}, \text{ (subject to a cap of 1)}. \tag{3.8}$$

Note that if the premium charged is sufficiently small, it is possible for a risk-group to be fully insured, i.e. $d_i(\pi_i) = 1$.

For iso-elastic demand, the equilibrium condition in Equation (3.4) takes the form:

$$E\left(\underline{\pi}\right) = \sum_{i=1}^{n} p_i\, \tau_i \left(\frac{\mu_i}{\pi_i}\right)^{\lambda_i} (\pi_i - \mu_i) = 0. \tag{3.9}$$

3.5. Loss coverage

In Section 1.1, loss coverage was defined as: expected losses compensated by insurance for the population as a whole, i.e.:

$$\text{Loss coverage: } C\left(\underline{\pi}\right) = \sum_{i=1}^{n} p_i\, d_i(\pi_i)\, \mu_i. \tag{3.10}$$

For iso-elastic demand, the expression for loss coverage takes the form:

$$C\left(\underline{\pi}\right) = \sum_{i=1}^{n} p_i\, \tau_i \left(\frac{\mu_i}{\pi_i}\right)^{\lambda_i} \mu_i. \tag{3.11}$$

We suggest that a good objective for a regulator or policymaker is to calibrate the price collar κ to maximise loss coverage over all possible premium regimes. This will ensure that voluntary purchases of insurance cover the largest possible fraction of the population's losses.[8] Mathematically, the objective can be stated, in terms of premiums, as[9]:

$$\max_{\kappa}\ C\left(\underline{\pi}\right), \text{ subject to } E\left(\underline{\pi}\right) = 0. \tag{3.12}$$

4. The case of two risk-groups

4.1. Maximising loss coverage

Consider two premium regimes: $\underline{\pi} = (\pi_1, \pi_2)$ and $\underline{\pi} + \underline{\Delta\pi} = (\pi_1 + \Delta\pi_1, \pi_2 + \Delta\pi_2)$, where both regimes satisfy the equilibrium condition in Equation (3.9), so that $E\left(\underline{\pi} + \underline{\Delta\pi}\right) = E\left(\underline{\pi}\right) = 0$. If $\underline{\Delta\pi}$ is "small", ignoring higher-order terms in the Taylor series expansion gives:

$$\Delta E = E\left(\underline{\pi} + \underline{\Delta\pi}\right) - E\left(\underline{\pi}\right) = E_1\, \Delta\pi_1 + E_2\, \Delta\pi_2, \text{ where } E_i = \frac{\partial E}{\partial \pi_i} \text{ for } i = 1, 2. \tag{4.1}$$

As $E\left(\underline{\pi} + \underline{\Delta\pi}\right) = E\left(\underline{\pi}\right) = 0$, and thus $\Delta E = 0$, the relationship between $\Delta\pi_1$ and $\Delta\pi_2$ can be expressed as:

$$\Delta\pi_2 = -\frac{E_1}{E_2}\, \Delta\pi_1. \tag{4.2}$$

[8] Compulsory purchase (a mandate) could ensure 100% coverage, but the reduction in liberty this involves is often unacceptable (this is discussed further in Section 6.1).

[9] Readers who are familiar with Lagrange multipliers and the Kuhn-Tucker theorem (please see Dixit (1990) for an exposition from an economic perspective) will realise that the constrained maximisation problem can be framed in terms of these optimisation approaches. However, instead of applying these methods mechanically, we provide the intermediate steps, so that the underlying economic interpretations are not overlooked.

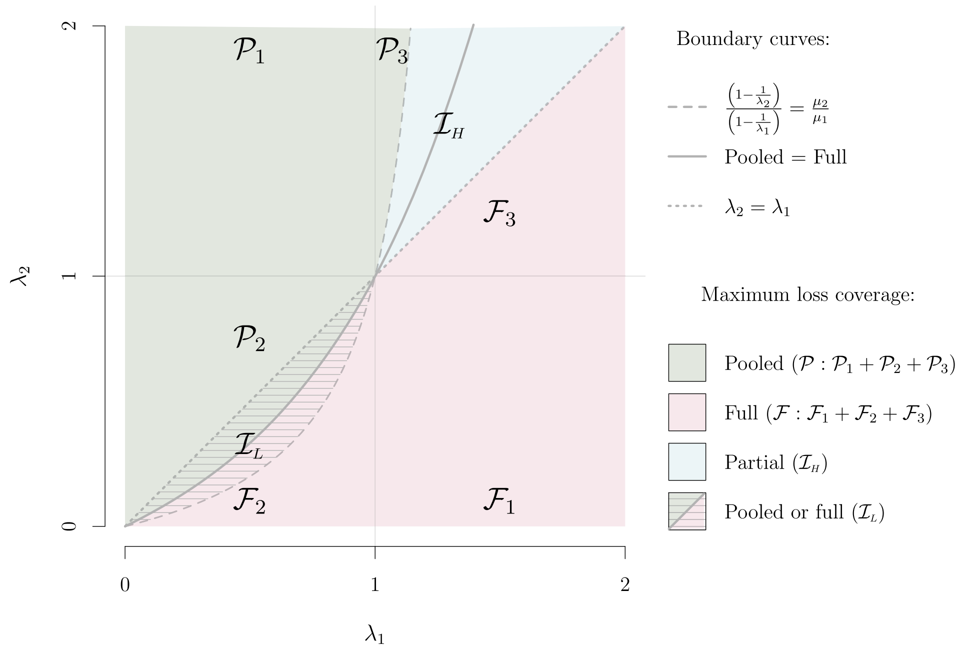

Fig. 3. Regions of (λ_1, λ_2)–plane where pooling, full or partial risk-classification give highest loss coverage. (Basis: $p_1 = p_2 = 1/2$, $\tau_1 = \tau_2 = 1/2$, $\mu_2/\mu_1 = 4$, but similar appearance for other plausible parameters.) (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

To compare loss coverages under two equilibrium premium regimes: $\underline{\pi} = (\pi_1, \pi_2)$ and $\underline{\pi} + \underline{\Delta\pi} = (\pi_1 + \Delta\pi_1, \pi_2 + \Delta\pi_2)$, Taylor series expansion ignoring higher-order terms gives:

$$\Delta C = C_1\,\Delta\pi_1 + C_2\,\Delta\pi_2 = \left[C_1 - \frac{E_1}{E_2}\,C_2\right]\,\Delta\pi_1, \text{ where } C_i = \frac{\partial C}{\partial \pi_i} \text{ for } i = 1, 2. \tag{4.3}$$

For the specific case of iso-elastic demand functions, after obtaining E_i and C_i by taking derivatives, the sensitivity of loss coverage to small changes in the equilibrium premium regimes is given by:

$$\Delta C = T\,\left(m_2 - m_1\right)\,\Delta\pi_1 \tag{4.4}$$

$$\text{where} \quad T = \frac{p_1\,\tau_1\,\lambda_1\,\left(\frac{\mu_1}{\pi_1}\right)^{\lambda_1+1}}{1 - m_2} \tag{4.5}$$

$$\text{and} \quad m_i = \frac{\pi_i}{\mu_i}\left(1 - \frac{1}{\lambda_i}\right) \quad \text{for } i = 1, 2. \tag{4.6}$$

Note that the term T is always positive, because the political constraint requires $\pi_2 \le \mu_2$, and thus $m_2 < 1$.

The sign of the term $\left(m_2 - m_1\right)$ determines how C depends on the low-risk premium π_1. Specifically, for given values of λ_1 and λ_2, if the term $\left(m_2 - m_1\right)$ is positive or negative over the whole politically feasible range of π_1, this tells us that loss coverage is maximised for pooling or full risk-classification respectively; and if it changes sign over the range of π_1, there will be a turning point for loss coverage. A further interpretation of the m_i will be given in Section 4.2 below.

The result is illustrated in Fig. 3. A formal statement as Theorem 2 and a proof are given in Appendix B.

Fig. 3 shows which risk-classification scheme – pooled, full or partial – gives highest loss coverage in each region of the (λ_1, λ_2)–plane. There are three main regions:

- $\mathcal{P}$ (left, green) where pooling is best.
- $\mathcal{F}$ (right, red) where full risk-classification is best.
- $\mathcal{I}_H$ (intermediate, high elasticities, blue) where for each point in the region, there is a particular partial risk-classification regime which gives an interior *maximum*.
- The hatching $\mathcal{I}_L$ (intermediate, low elasticities) over parts of the green and red regions denotes a zone where for each point, a particular partial risk-classification regime gives an interior *minimum* (note the analogy with region $\mathcal{I}_H$, where partial gives an interior *maximum*). This is consistent with the underlying green and red coding, i.e. one of pooling and full risk-classification regimes must be the maximum.

Also note:

- The large green region (pooling best) and red region (full best) are each divided into three sub-regions, delineated by the vertical and horizontal lines $\lambda_1 = 1$ and $\lambda_2 = 1$.
- The grey central curve along which pooling and full risk-classification give equal loss coverage is the same as previously shown in Fig. 1.

Table 1
Estimates of demand elasticity for various insurance markets.

Market and country	Demand elasticities[a]	Authors
Health insurance exchange, USA	2.05 (ages 26-31), 1.36 (ages 62-64)	Tebaldi (2022)
Health insurance exchange, USA	young > 2× old	Ericson and Starc (2015)
Health insurance (pre-ACA), USA	0 to 0.2	Chernew et al. (1997), Blumberg et al. (2001), Buchmueller and Ohri (2006)
Health insurance, Australia	0.35 to 0.50	Butler (1999)
Long-term care insurance, USA	3.3	Goda (2011)
Term life insurance, USA	0.66	Viswanathan et al. (2006)
Yearly renewable term life, USA	0.4 to 0.5	Pauly et al. (2003)
Whole life insurance, USA	0.71 to 0.92	Babbel (1985)
Farm crop insurance, USA	0.32 to 0.73	Goodwin (1993)

[a] Estimates in empirical papers are generally given as negative values, but we have presented the absolute values here for consistency with the definition of demand elasticity used in this paper.

The range of elasticities in Fig. 3 is intended to be representative of real-world product demand elasticities.[10] In this regard, Table 1 shows estimates from various authors for a selection of insurance markets. Note particularly Tebaldi (2022) and Ericson and Starc (2015), who provide some evidence of different elasticities for different risk-groups: they report substantial differences in elasticity between younger and older customers (i.e. lower and higher risks) on health insurance exchanges.

For graphical intuition into the general pattern of Fig. 3, look at Fig. 4, where the six panels show representative plots of equilibrium loss coverage against π_1 for all the labelled regions in Fig. 3. In each panel, increasing π_1 along the x-axis corresponds to reducing the price collar from its maximum of $\kappa = \mu_2/\mu_1$ (which permits full risk-classification) to its minimum of $\kappa = 1$ (which enforces pooling). Then note that the slope and convexity of the various plots account for the pattern of results illustrated in Fig. 4. In particular:

(i) For pooling to be best, the plot needs to have a maximum at its right extreme: regions $\mathcal{P}_1$, $\mathcal{P}_2$ & $\mathcal{P}_3$ (first row of Fig. 4).
(ii) For full risk-classification to be best, the plot needs to have a maximum at its left extreme: regions $\mathcal{F}_1$, $\mathcal{F}_2$ & $\mathcal{F}_3$ (second row of Fig. 4).
(iii) For partial risk-classification to be best, the plot needs to have an interior maximum (right panel of third row of Fig. 4). This is true only in the upper-right intermediate region $\mathcal{I}_H$.
(iv) For partial risk-classification to be worst, the plot needs to have an interior minimum (left panel of third row of Fig. 4). This is true only in the lower-left intermediate region $\mathcal{I}_L$.

4.2. Economic intuition: marginal revenue-to-cost ratios

From Equation (4.4), $\Delta C/\Delta\pi_1 = T\,(m_2 - m_1)$, we can see that the slope of the loss coverage plots in Fig. 4 depends on $(m_2 - m_1)$. Each m_i represents the ratio of *marginal revenue* to *marginal cost* for the risk-group.

To see this, first note that the total revenue and total cost given risk-group i are $TR_i(\pi_i) = \pi_i\, d_i(\pi_i)$ and $TC_i(\pi_i) = \mu_i\, d_i(\pi_i)$. So the marginal revenue and marginal cost are:

$$MR_i(\pi_i) = \frac{\partial TR_i(\pi_i)}{\partial d_i(\pi_i)} = \frac{\partial\left[\pi_i\, d_i(\pi_i)\right]}{\partial d_i(\pi_i)} = \pi_i + d_i(\pi_i)\frac{\partial \pi_i}{\partial d_i(\pi_i)} = \pi_i\left(1 - \frac{1}{\epsilon_i(\pi_i)}\right) \tag{4.7}$$

$$MC_i(\pi_i) = \frac{\partial TC_i(\pi_i)}{\partial d_i(\pi_i)} = \mu_i. \tag{4.8}$$

So we define the *marginal revenue-to-cost ratio* as:

$$m_i = \frac{MR_i(\pi_i)}{MC_i(\pi_i)} = \frac{\pi_i}{\mu_i}\left(1 - \frac{1}{\lambda_i}\right) \text{ for iso-elastic demand,} \tag{4.9}$$

which is the same as the m_i previously identified in Equation (4.6). The ratio measures the extent to which a small increase in a risk-group's cost (i.e. its demand, multiplied by its risk rate μ) is offset by a corresponding increase in its revenue (i.e. its demand, multiplied by its premium rate π).

The entirety of Fig. 3 can then be characterised by the following principle: *to increase loss coverage, shift demand towards the risk-group with higher marginal revenue-to-cost ratio.*[11] This seems very intuitive: it helps to shift demand towards the risk-group which is more effective in generating incremental revenue to cover its incremental cost.

The principle is illustrated in Fig. 5. Throughout the solid green region, the ratio is higher for the high risk-group, for all feasible values of the collar κ (i.e. all premiums π_1 and π_2). So it helps to shift demand maximally towards the high risk-group (i.e. pooling is best). The converse applies throughout the solid red region.

In the blue region, either ratio can be higher, depending on the current value of the collar κ. So it helps to shift demand "so far, but no further", until the two ratios are equal (i.e. neither is higher), where loss coverage reaches an interior maximum.

In the hatched zone, either ratio can again be higher, depending on the current collar κ; but when they are equal, loss coverage now has an interior minimum. From this point, either increasing or decreasing the collar κ will increase loss coverage; and so to be sure of maximising, we need

[10] *Product* elasticities, rather than *brand* elasticities, are relevant to our analysis (cf. Footnote 6).
[11] "Higher" here means either more positive, or less negative (if both ratios are negative, i.e. both demands are inelastic). For proof of the principle see Lemma 3 in Appendix B.

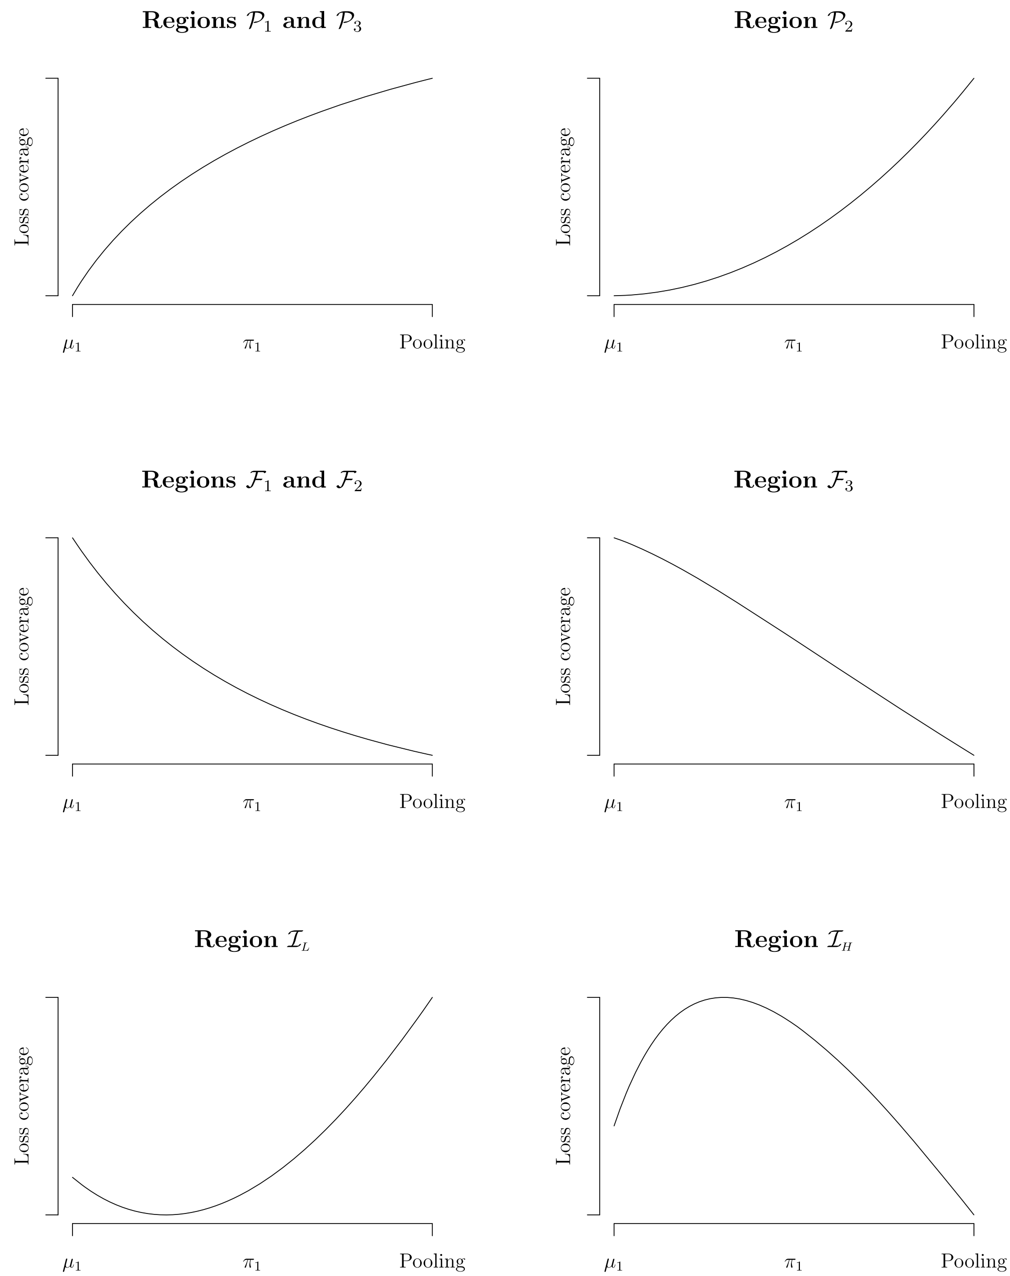

Fig. 4. Slope and convexity of loss coverage–premium plots.

to check which extremum (i.e. pooling or full risk-classification) gives the largest increase. (To see this graphically, look at the bottom left panel in Fig. 4. If we start to the left of the minimum point, marginally increasing loss coverage will take us towards the left extremum, but the right extremum gives the highest loss coverage.)

To see why the turning point in the blue region is a maximum, first note that the "equal ratios" condition for a turning point is:

$$m_1 = m_2 \quad \Rightarrow \quad \frac{\pi_1}{\mu_1}\left(1 - \frac{1}{\lambda_1}\right) = \frac{\pi_2}{\mu_2}\left(1 - \frac{1}{\lambda_2}\right). \tag{4.10}$$

Now note that in the blue region, both $\lambda_i > 1$, so both brackets above are always positive. So starting from the turning point, if we move $\pi_1 \downarrow$ and $\pi_2 \uparrow$ (i.e. shift demand towards low risks, and away from high risks), the marginal revenue-to-cost ratio *falls* in the risk-group with *rising* demand, and vice versa. This is an unfavourable combination. The same unfavourable combination obtains if we move $\pi_1 \uparrow$ and $\pi_2 \downarrow$. So whichever way we move the premiums to give a new equilibrium, loss coverage decreases; and therefore the turning point we started from must be a maximum.

Conversely, in the hatched zone, both $\lambda_i < 1$, and so both brackets above are always negative. The entire argument just given is then inverted, so the turning point for loss coverage in the hatched zone must be a minimum.

4.3. A general rule to select the optimal collar

The required price collar to equalise the marginal revenue-to-cost ratios and so give a turning point as just discussed can be found from Equation (4.10) as:

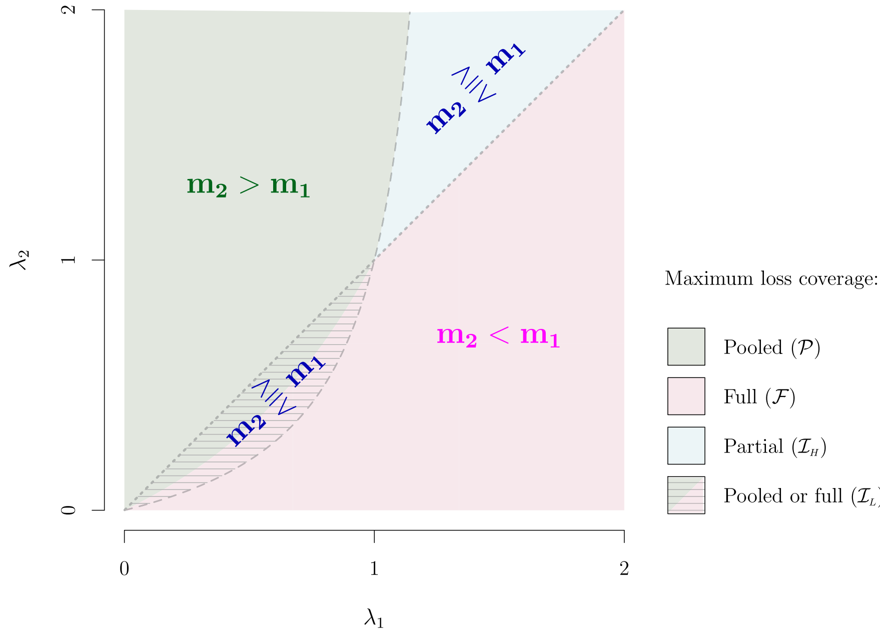

Fig. 5. Relationship between the marginal revenue-to-cost ratios m_1 and m_2 in the different regions of the (λ_1, λ_2)–plane.

$$\kappa = \frac{\pi_H}{\pi_L} = \frac{\pi_2}{\pi_1} = \frac{\frac{1}{\mu_1}\left(1-\frac{1}{\lambda_1}\right)}{\frac{1}{\mu_2}\left(1-\frac{1}{\lambda_2}\right)} = \Psi \text{ say,} \tag{4.11}$$

but in regions other than $\mathcal{I}_L$ and $\mathcal{I}_H$, this turning point will lie outside the feasible range $[1, \mu_2/\mu_1]$ for the collar.[12]

When we limit the collar to the feasible range, the following rule selects the optimal collar κ to maximise loss coverage:

- for $\lambda_2 > 1$ set κ as *close* as feasible to Ψ and
- for $\lambda_2 < 1$: set κ as *distant* as feasible from Ψ.

In region $\mathcal{I}_H$, this rule sets $\kappa = \Psi$ to give the turning point (a maximum) as in Equation (4.11). In all other regions, it selects the feasible extremum for κ (i.e. pooling or full risk-classification) which maximally implements the principle: *to increase loss coverage, shift demand towards the risk-group with higher marginal revenue-to-cost ratio.*

The form of the rule and its dependence on $\lambda_2 \lesseqgtr 1$ can also be understood graphically by reference to the curvature apparent in the six panels in Fig. 4. In all panels where $\lambda_2 > 1$, the curve is concave, so the theoretical turning point given by Equation (4.11) (not necessarily politically feasible) is a maximum. For these panels, the rule tells us to select the collar which gets us as close as possible to that turning point; and vice versa in all panels where $\lambda_2 < 1$, where the curve is convex and so the turning point is a minimum. (Although the theoretical turning point is not visible in the top four panels, its "off-screen" location is always visually implied: the flatter end of the curve points towards the turning point.)

5. Extensions and sensitivities

5.1. More than two risk-groups

For more than two risk-groups, we can generalise Theorem 2, using:

(i) the known tripartite structure of the Nash equilibrium premium regime under a price collar, as illustrated in Fig. 2 and formalised in Theorem 1, and
(ii) a heuristic assumption that all risk-groups in the "low" super-group $\mathcal{L}$ have the same iso-elastic demand elasticity λ_L, and all risk-groups in the "high" super-group $\mathcal{H}$ have the same iso-elastic demand elasticity λ_H. This is based on the premise that risk-groups with broadly similar risks are likely to have broadly similar elasticities.

Intuitively, the key insight for more than two risk-groups is that because all risk-groups in the "middle" super-group $\mathcal{M}$ always pay their fair premiums, they do not contribute to the cross-subsidies which determine equilibrium (and hence loss coverage). So for the purpose of determining how loss coverage changes when the collar changes, the "middle" super-group $\mathcal{M}$ can be completely disregarded. Therefore the analysis of loss coverage for two risk-groups in Equation (4.4) can be re-stated in terms of the two super-groups $\mathcal{L}$ and $\mathcal{H}$, with the super-group elasticities λ_L and λ_H in place of λ_1 and λ_2, and other parameters also set to their super-group values, i.e.:

[12] To illustrate, in region $\mathcal{P}_2$, there is a turning point with $\pi_1 < \mu_1$ and $\pi_2 > \mu_2$, i.e. hyper-differentiated premiums, which are politically infeasible. In region $\mathcal{P}_1$, both revenues always increase on any move towards pooling, so the "technical" turning point must involve one *negative* premium. Continuity with hyper-differentiation in $\mathcal{P}_2$ as just discussed suggests that the negative one should be π_1; but any negative premium is both politically and economically infeasible.

$$\Delta C = T \left(m_H - m_L\right) \Delta\pi_L, \tag{5.1}$$

$$\text{where } T = \frac{p_L\, \tau_L\, \lambda_L \left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L+1}}{1 - m_H}; \text{ and } m_i = \frac{\pi_i}{\mu_i}\left(1 - \frac{1}{\lambda_i}\right); \quad i = L, H; \tag{5.2}$$

where for the "low" and "high" super-groups $\mathcal{L}$ and $\mathcal{H}$:

- μ_L and μ_H are the respective pooled equilibrium premiums;
- π_L and π_H are the corresponding premiums;
- p_L and p_H are the proportions of the population belonging to the super-groups;
- τ_L and τ_H can be interpreted as the 'fair-premium demand' when all risk-groups in the respective super-groups are pooled and charged the same pooled premium.

As T is always positive, the same inferences as before can then be made. A formal derivation of Equation (5.1) is given in Appendix C. Our general principle: *to increase loss coverage, shift demand towards the risk-group with higher marginal revenue-to-cost ratio* also remains valid when applied to the super-groups $\mathcal{L}$ and $\mathcal{H}$.

It is possible that when the width κ of the collar changes, the constituent risk-groups in the super-groups $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{H}$ may slightly change. This is not problematic, for the following reasons:

- Starting from pooling (i.e. $\kappa = 1$), all risk-groups with risks less than the pooled premium belong to the super-group $\mathcal{L}$ having the same demand elasticity λ_L; while the remaining risk-groups with higher risks belong to the super-group $\mathcal{H}$ having the same demand elasticity λ_H.
- As the price collar, κ, is increased, more and more risk-groups from the super-groups $\mathcal{L}$ and $\mathcal{H}$ join $\mathcal{M}$. However, the remaining risk-groups in $\mathcal{L}$ and $\mathcal{H}$ continue to have the same demand elasticities λ_L and λ_H respectively. So, any change in the compositions of $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{H}$ would not affect the underlying demand elasticities of the super-groups $\mathcal{L}$ and $\mathcal{H}$.
- Risk-groups at the upper end of $\mathcal{L}$ or lower end of $\mathcal{H}$, which are on the threshold of moving to $\mathcal{M}$, are *already* paying very close to their fair premiums. So their contribution to loss coverage is almost unchanged when they cross the threshold and transfer into $\mathcal{M}$.
- So provided insurance demands are "well-behaved" functions of premiums, which is indeed the case for iso-elastic demand functions, the loss coverage would be a continuous function of the price collar. Hence the conclusions would remain unaffected if the constituents of $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{H}$ change when the price collar changes.

5.2. Other demand functions

To apply the results for non-iso-elastic demand functions, we can use the construct of *arc elasticity* (Hao et al. (2018), Vazquez (1995)) between two prices a and b, defined as:

$$\text{arc elasticity: } \eta_i(a,b) = \frac{\int_a^b \epsilon_i(s)\, d\log s}{\int_a^b d\log s}. \tag{5.3}$$

Arc elasticity can be interpreted as the *average* value of (point) demand elasticity over the price logarithmic arc from price a to price b. The political constraints on prices imply feasible arcs for π_i to be (μ_i, π_0) or (π_0, μ_i) depending on whether or not μ_i is less than π_0. If we re-define λ_i as the corresponding arc elasticity, i.e.:

$$\lambda_i = \eta_i(\mu_i, \pi_i) = \frac{\int_{\pi_i}^{\mu_i} \epsilon_i(s)\, d\log s}{\int_{\pi_i}^{\mu_i} d\log s}, \tag{5.4}$$

we can then apply the techniques presented in previous sections using these λ_i-s, to compare various price collar regimes.

Note also that if we are evaluating only a small change in the risk-classification scheme (i.e. a small change in the high-risk and low-risk premiums), then only short arcs encompassing the proposed premium changes need to be considered. Furthermore, if elasticity is either less than 1 or greater than 1 for a risk-group throughout its relevant arc, then so is the arc elasticity. This allows quick inferences to be made without calculating arc elasticities, if one is prepared to assume either "elastic demand" or "inelastic demand" for a risk-group throughout its relevant arc.

5.3. Different population structures

By "population structure", we mean the proportions of low and high risks in the population, p_1 and p_2. These parameters were included in our definition of equilibrium in Equation (3.9), but do not appear in the definitions of most of the boundaries in Fig. 3. So making the population structure more extreme has no effect on most of the boundaries. The only change is that the central grey line of equality (i.e. where pooled and full give the same loss coverage) flexes slightly, while still passing through the points $(0,0)$ and $(1,1)$ and always remaining within the intermediate regions $\mathcal{I}_L$ and $\mathcal{I}_H$. This is illustrated for $p_2/p_1 = 50/50$ and 99/1 in the first panel in Fig. 6.[13]

Making the population structure more extreme also makes all the loss coverage plots in Fig. 4 flatter over their full range. In other words, if either risk-group is very small, it makes almost no difference to loss coverage which risk-classification scheme we use.

[13] For region $\mathcal{I}_L$, the exact position of the grey line of equality is given by Corollary 2.1 in Appendix B. A similar derivation could be given for region $\mathcal{I}_H$, but it is of little interest, because partial risk-classification gives the highest loss coverage in this region.

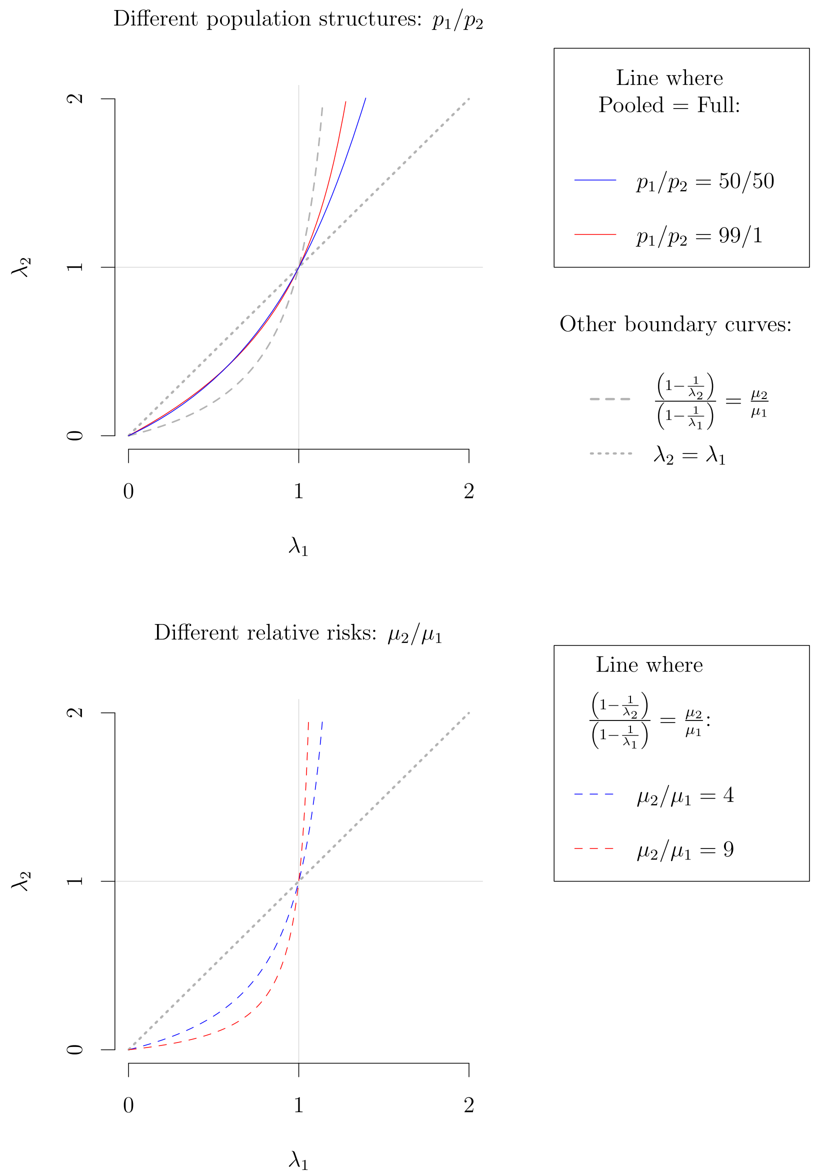

Fig. 6. Sensitivity to different population structures and relative risks.

5.4. Different relative risks

By "relative risk", we mean the ratio of high-risk to low-risk probabilities of loss, μ_2/μ_1. The main effect of higher μ_2/μ_1 is to shrink the lower right region $\mathcal{F}_2$ inside the unit square where full risk-classification is always best. This effect is shown in the second panel of Fig. 6. The intuition is that as $\mu_2/\mu_1 \to \infty$, the relative contribution that low risks can potentially make to loss coverage falls. This implies there is less benefit in increasing demand from the low risks (i.e. reducing their premiums, via full risk-classification). So the range of elasticity combinations for which full risk-classification is best shrinks.

Comparing the two panels in Fig. 6, the effect of a modest change in relative risk is more noticeable than the effect of an extreme change in population structure. This is because changing relative risk shifts only costs between the risk-groups, but changing population structure shifts both costs and premiums, which to a large degree offset one another. The former is more disruptive of any initial equilibrium.

6. Discussion

6.1. Loss coverage versus other policy metrics

We have assumed that the objective – or at least, one objective – of insurance policymakers is to promote compensation of the population's losses. This objective motivates the concept of loss coverage (expected losses covered), with risk-classification schemes which generate higher loss coverage being preferred. But the concept is not widely used in policy analysis, which prompts the questions: do policymakers actually care about loss coverage; and if not, what do they care about instead?

The most definitive way of promoting loss coverage is to make insurance compulsory. The fact that this is done in many jurisdictions, for some important classes of insurance, suggests that policymakers do implicitly care very much about loss coverage. Compulsion as an expedient to maximise loss coverage is commonly applied where any uninsured losses would fall on innocent third parties (e.g. auto and employer liability insurances), or where they relate to universal needs (e.g. some basic level of healthcare).

In other classes of personal insurance, such as life insurance, the third-party or universal-need justifications for compulsion are less clear-cut. On the one hand, people with no dependants have little or no need for life insurance. On the other hand, people who do have dependants often have far too little life insurance (Auerbach and Kotlikoff (1991), Bernheim et al. (2003)). In these circumstances, compulsion may be a poorly targeted policy for increasing loss coverage. It may be more politically acceptable to leave purchasing decisions to individuals, and look to adjust risk classification (e.g. the width of price collar) as a “nudge” to increase loss coverage.

Whilst increasing loss coverage seems an implicit motivation for many compulsory insurances, most political discussion focuses on coverage of persons rather than coverage of losses. Our impression is that this is not based on any principled preference, bur rather on the heuristic that since coverage of persons must always precede coverage of losses, the two are equivalent and no distinction needs to be made (but we say that it does). Much applied policy analysis, perhaps taking its cue from the political discussion, also focuses on coverage of persons. As recent examples, Wettstein (2017), Gruber and Sommers (2019) and Einav et al. (2019) all focus mainly on increases in enrolment (i.e. coverage of persons) under the Affordable Care Act.

More theoretical policy analyses, typically directed at academic rather than policy audiences, focus on utilitarian social welfare. One approach makes assumptions about utility functions and then directly calculates expected utility for a random member of the population, behind a Rawlsian veil of ignorance that screens off knowledge of one’s risk type (Hoy (2006)). Alternatively, on the assumption of revealed preference (i.e. demand from high and low risks represents their willingness-to-pay and hence expected utility), observed movements in the demand curves can be used to estimate changes in consumer and social surplus as measures of changes in welfare (Einav et al. (2010)). The latter approach has become popular in empirical work in recent years (for a survey of applications see Einav and Finkelstein (2023)), but we think it has two main drawbacks.

First, the principle of revealed preference assumes that insurance decisions accurately reflect expected utilities, but much evidence calls this into question. Probability distortions appear substantial, particularly over-weighting of the small probabilities usually relevant for insurance (Barseghyan et al. (2013)); inertia, lack of information and time and hassle costs appear to greatly influence decisions (Handel (2013), Ericson (2014), Handel and Kolstad (2015)); and more than half of health plan participants in one major study actually chose financially dominated options (Bhargava et al. (2017)). One justification for a simple focus on expected losses compensated (i.e. loss coverage) is that it sidesteps all such distortions.

Second, the demand curve approach implicitly places the same weights on changes in demand from high and low risks. However, footnote 6 in Einav and Finkelstein (2023) notes (but does not pursue) the important point that since higher risks are worse off *ex-ante*, this should imply higher weights for them in any social welfare function evaluated behind the veil of ignorance. Similarly, Hendren (2021) notes that willingness-to-pay is observed at a time when individuals’ risk status is already known, and so misses the value of insuring against high-risk status *ex ante* behind the veil of ignorance. Both these points are “directionally” similar to loss coverage; we say that it is sensible to place higher weight *on* demand from higher risks, in linear proportion to their higher expected losses.[14]

A practical advantage of loss coverage compared with welfare is that its evaluation requires knowledge only of numbers insured in each risk-group and their expected losses, rather than unobservable utility functions or implicit links between utility functions and demand. This simplicity makes loss coverage more amenable to observation and measurement, and also to explanation and advocacy to broad policy audiences.

6.2. Price collars versus bans on specific variables

A common regime of risk-classification in retail insurance markets is “full, but with bans on a few deprecated variables (e.g. gender, age, genetic tests, etc)”. Conceptually, this “full with restrictions” risk-classification is not quite the same as our price collar. But if insurers are unable to fully compensate for the banned variables by using correlated variables, “full with restrictions” seems likely to lead to a similar outcome to a price collar: low risks are charged a bit more than their fair premiums, and high risks a bit less. Our result in Fig. 3 then seems rather striking: partial risk-classification can never be optimal for loss coverage if both demand elasticities (or arc elasticities over the relevant arcs) are less than 1.

However, this result says nothing about the quantum of the difference in loss coverage between any particular partial regime and the optimal regime (be it full or partial risk-classification). Partial risk-classification is not a single third option, but rather a continuum of options between full risk-classification and pooling. If full risk-classification is optimal, a “broad” price collar (or a ban on a variable with limited relevance) enforces a partial regime very close to full risk-classification. Conversely, if pooling is optimal, a “narrow” price collar enforces a partial regime very close to pooling. In either case, the reduction in loss coverage compared to the optimum will be very small. Policymakers may regard this as an acceptable cost for the wider social benefits of restrictions on deprecated variables.

An advantage of a price collar calibrated to maximise loss coverage, compared with restrictions on particular variables, is that it does not privilege any particular risk-groups, or any particular causes of increased risk. It simply places higher weights on all higher risks, but only in proportion to their higher expected losses. In this sense, it side-steps critiques along the lines of “genetic exceptionalism” (e.g. Mittra (2006), Malpas (2008), Murray (2019)) sometimes directed at bans on specific variables such as genetic test results (i.e. the critique that it is unjustified to ban these variables, when many other variables have similar effects on risk).

6.3. Price collar versus price cap

Any equilibrium implemented by a price collar can alternatively and equivalently be implemented by a price cap, that is a single upper limit π_H on the premiums insurers can charge.[15] The proof of the equivalence is given in Corollary 1.1 in Appendix A.

However, maximising loss coverage via a price cap has the technical disadvantage that we need to find the π_H that solves Equation (3.12), i.e. maximising loss coverage subject to satisfying the equilibrium condition. This requires knowledge of all of the risks (μ_i), demand elasticities (λ_i), population proportions (p_i) and fair-premiums demands (τ_i).

On the other hand, the optimal price collar from the feasible range $[1, \mu_2/\mu_1]$ is given by the rule in Section 4.3:

[14] Linearity places equal weights on compensating expected losses of high and low risks. There might sometimes be policy arguments to modify this weighting scheme, e.g. to reflect advantages or disadvantages associated with particular risk-groups that are not captured by their risk status. But for a policymaker who cares about compensation of losses, “equal weights on equal expected losses” seems the obvious default.

[15] See Dosis (2022) for an investigation into regulatory price caps with endogenous contracts in the tradition of Rothschild and Stiglitz (1976) and Wilson (1977).

- for $\lambda_2 > 1$ set κ as close as feasible to Ψ; and
- for $\lambda_2 < 1$: set κ as distant as feasible from Ψ,

and this requires knowledge only of the terms in Ψ, i.e. the μ_i and λ_i.

Apart from the technical disadvantage of greater information requirements, a price cap also has potential political disadvantages. While a correctly calibrated price cap can implement any price collar, it has the appearance of a more severe intervention: a constraint on the overall level of prices rather than just their spread. Also, given our necessary assumption of guaranteed issue (sub-section 2.2), a badly calibrated price cap that is set too low may not allow a zero-profit equilibrium at all.

7. Conclusions

Loss coverage, defined as expected population losses compensated by insurance, is a public policy criterion for comparing different risk-classification regimes. Using a model with two risk-groups (high and low) and iso-elastic demand, we compared loss coverage under three alternative regulatory regimes: (i) full risk-classification (ii) pooling (iii) a price collar, whereby each insurer is permitted to set any premiums, subject to a maximum ratio of its highest and lowest prices for different risks. Outcomes depend on the comparative demand elasticities of low and high risks. If low-risk elasticity is sufficiently low compared with high-risk elasticity, pooling is optimal; and if it is sufficiently high, full risk-classification is optimal. For an intermediate region where the elasticities are not too far apart, a price collar is optimal, but only if both elasticities are greater than one.

A key driver of these results is the ratio of marginal revenue to marginal cost for each risk-group. From any initial equilibrium, the general principle is: *to increase loss coverage, shift demand towards the risk-group with higher marginal revenue-to-cost ratio*.

For more than two risk-groups, the results can be extended via the insight that equilibrium always involves a tripartite arrangement of risk-groups into three super-groups: a super-group $\mathcal{L}$ of low risk-groups all charged the same low premium, a super-group $\mathcal{H}$ of high risk-groups all charged the same high premium, and a middle super-group $\mathcal{M}$ all charged their actuarially fair premiums.

For non-iso-elastic demand functions, the results can be extended using the construct of arc elasticity of demand, which can be thought of as the average value of point elasticity over the logarithmic arc between two prices.

Declaration of competing interest

None.

Data availability

No data was used for the research described in the article.

Acknowledgement

We thank Angus Macdonald for useful comments. Indradeb Chatterjee and MingJie Hao thank Radfall Charitable Trust for research funding.

Appendix A. Theorem 1

Theorem 1. *If there are n risk-groups, with risks $\mu_1 < \mu_2 < \cdots < \mu_n$ with a price collar of κ, where $1 \le \kappa \le \mu_n/\mu_1$, there exists a Nash equilibrium premium regime $\underline{\pi} = (\pi_1, \pi_2, \ldots, \pi_n)$, such that:*

$$\pi_i = \begin{cases} \pi_L & \text{if } \mu_i < \pi_L; \\ \mu_i & \text{if } \pi_L \le \mu_i \le \pi_H; \\ \pi_H & \text{if } \mu_i > \pi_H, \end{cases} \tag{A.1}$$

where $\pi_L = \min_i \pi_i$, $\pi_H = \max_i \pi_i$ and $\pi_H = \kappa\,\pi_L$.

Proof. We will prove the theorem using the following steps:

1. An equilibrium premium regime with the structure proposed in Equation (A.1) exists.
2. If there are multiple equilibrium premium regimes with the same proposed structure, the regime with the smallest π_L is preferred among all such regimes.
3. Given π_L and $\pi_H = \kappa\,\pi_L$, the premium regime with the proposed structure cannot be destabilised by any other equilibrium premium regime with the same π_L and π_H but having a different structure.
4. Given a compulsory price collar κ, the premium regime with the proposed structure cannot be destabilised by any other premium regime based on a voluntary smaller price collar.

Proof of step 1. Given a price collar κ, define the expected profit from setting the lowest premium π_L, where $\mu_1 \le \pi_L \le \mu_n/\kappa$, as follows:

$$e_\kappa\left(\pi_L\right) = E\left(\underline{\pi}\right) = \sum_{i=1}^{n} p_i\, d_i(\pi_i)\, \left(\pi_i - \mu_i\right); \text{ where } \pi_i = \begin{cases} \pi_L & \text{if } \mu_i < \pi_L; \\ \mu_i & \text{if } \pi_L \le \mu_i \le \kappa\,\pi_L; \\ \kappa\,\pi_L & \text{if } \mu_i > \kappa\,\pi_L. \end{cases} \tag{A.2}$$

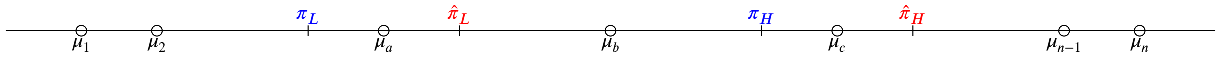

Fig. 7. Theorem 1: Proof of Step 2.

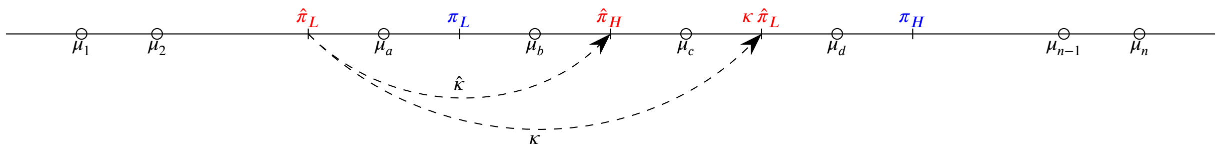

Fig. 8. Theorem 1: Proof of Step 4.

If $\pi_L = \mu_1$, as $\kappa\,\mu_1 \leq \mu_n$, expected profit cannot be positive, i.e.: $e_\kappa\left(\mu_1\right) \leq 0$.

If $\pi_L = \mu_n/\kappa$, as $\mu_1 \leq \mu_n/\kappa$, expected profit cannot be negative, i.e.: $e_\kappa\left(\mu_n/\kappa\right) \geq 0$.

Assuming continuity of the demand functions $d_i(\pi_i)$ for all risk-groups, $e_\kappa(x)$ is also a continuous function. So, by the intermediate value theorem, there exists a value π_L, such that $\mu_1 \leq \pi_L \leq \mu_n/\kappa$, for which $e_\kappa\left(\pi_L\right) = 0$. This proves the existence of an equilibrium premium regime as outlined in the theorem. □

Proof of step 2. If there are multiple solutions to the equation, $e_\kappa\left(\pi_L\right) = 0$, the premium regime based on the smallest of these roots cannot be destabilised by premium regimes based on any other solutions of $e_\kappa\left(\pi_L\right) = 0$. To show this, suppose if possible there are two premium regimes:

$$\underline{\pi} = \left(\pi_1, \pi_2, \ldots, \pi_n\right), \text{ with } \pi_H = \kappa\,\pi_L, \text{ where } \pi_L = \min_i \pi_i \text{ and } \pi_H = \max_i \pi_i;$$

$$\underline{\hat{\pi}} = \left(\hat{\pi}_1, \hat{\pi}_2, \ldots, \hat{\pi}_n\right), \text{ with } \hat{\pi}_H = \kappa\,\hat{\pi}_L, \text{ where } \hat{\pi}_L = \min_i \hat{\pi}_i \text{ and } \hat{\pi}_H = \max_i \hat{\pi}_i;$$

with $\pi_L < \hat{\pi}_L$ (and consequently $\pi_H < \hat{\pi}_H$), such as shown in Fig. 7.

- $\underline{\hat{\pi}}$ involves higher premiums for all risk-groups with $\mu_i < \hat{\pi}_L$, and higher premiums for all risk-groups with $\mu_i > \pi_H$, and the same (actuarially fair) premiums to all risk-groups in between.
- So all risks are either not attracted to $\underline{\hat{\pi}}$ (because it charges more), or are indifferent (because it charges the same); and in the latter case, the premium regime $\underline{\pi}$ would also be indifferent to the loss of these risk-groups, because they generate no profit or loss under $\underline{\pi}$.

Therefore any higher solution $e_\kappa\left(\pi_L\right) = 0$ cannot destabilise the lowest solution. □

Proof of step 3. Any alternative regime $\underline{\hat{\pi}}$ with the same π_L and π_H and a different structure needs to charge some risks more, and/or some risks less.

- Risks which are charged less will be attracted by $\underline{\hat{\pi}}$, and generate a smaller profit contribution (or larger loss) than they do under $\underline{\pi}$.
- To achieve zero profit then requires that this deficit be made up by charging other risks more.
- But no risks will be prepared to pay more, because they will prefer $\underline{\pi}$.

Therefore $\underline{\pi}$ cannot be destabilised by an alternative regime $\underline{\hat{\pi}}$ with the same π_L and π_H and a different structure. □

Proof of step 4. As a regulatory price collar would only require $\pi_H \leq \kappa\,\pi_L$, we need to show that $\underline{\pi}$ cannot be destabilised by an equilibrium premium regime using a smaller price collar.

Consider an alternative equilibrium premium regime: $\underline{\hat{\pi}} = \left(\hat{\pi}_1, \hat{\pi}_2, \ldots, \hat{\pi}_n\right)$, with minimum and maximum premiums: $\hat{\pi}_L = \min_i \hat{\pi}_i$ and $\hat{\pi}_H = \max_i \hat{\pi}_i$ respectively, such that $\hat{\pi}_H = \hat{\kappa}\,\hat{\pi}_L$, for some $\hat{\kappa} < \kappa$, such as shown in Fig. 8.

- If the lower end of the new collar is set *above* the lower end of the old collar, then $\underline{\hat{\pi}}$ attracts none of the profitable low risks, and so cannot be profitable and cannot destabilise $\underline{\pi}$.
- If the lower end of the new collar is set *below* the lower end of the old collar, then $\underline{\hat{\pi}}$ generates a smaller profit contribution on all risk-groups with $\mu_i < \pi_L$.
- Equilibrium requires that this deficit be made up by charging other risks more.
- But because of the smaller $\hat{\kappa}$ and lower base of the collar, $\underline{\hat{\pi}}$ does not charge any risk-groups more.

Therefore no equilibrium solution with a smaller collar $\hat{\kappa}$ exists. □

So, $\underline{\pi}$, as outlined in the theorem, exists, and is a unique Nash equilibrium premium regime satisfying all political, regulatory and economic constraints. ■

Corollary 1.1. *A regulatory price cap regime, which specifies an upper limit, π_H, on the premiums that insurers can charge, produces the same stable equilibrium as under a price collar regime, given in Equation (A.1), as long as the value of π_H allows for viable equilibrium satisfying the political constraint.*

Proof. Assuming $\pi_H < \mu_n$, (as otherwise, it would lead to a full risk-classification regime) a price cap regime implicitly creates a super-group $\mathcal{H}$ of high risk-groups, all charged the same premium π_H. Then, under perfect competition, to ensure zero-profits equilibrium, insurers would have to recoup the losses from the remaining risk-groups. By the intermediate value theorem, there exists a $\pi_L \le \pi_H$, such that we get a premium regime $\underline{\pi}$ with the same tripartite pattern as in Equation (A.1) and depicted in Fig. 2, which satisfies the equilibrium condition.

This equilibrium premium regime, $\underline{\pi}$, cannot be destabilised by any other price cap premium regime $\hat{\underline{\pi}}$, because risks which are charged less will be attracted by $\hat{\underline{\pi}}$, and generate a smaller profit contribution (or larger loss) than they do under $\underline{\pi}$. To achieve zero profit then requires that this deficit be made up by charging other risks more. But no risks will be prepared to pay more, because they will prefer $\underline{\pi}$. So, $\underline{\pi}$ cannot be destabilised. □

Appendix B. Theorem 2

In order to prove Theorem 2, we first provide the algebraic inequalities defining the regions represented graphically in Fig. 3. Recall the definitions of graph regions: $\mathcal{P}$: "Pooled", $\mathcal{F}$: "Full risk-classification", $\mathcal{I}_L$: "Intermediate, low elasticity" and $\mathcal{I}_H$: "Intermediate, high elasticity", so that:

$\mathcal{I}_L$: $\lambda_2 \le \lambda_1 \le 1$ and $1 \le \left(1-\frac{1}{\lambda_2}\right)/\left(1-\frac{1}{\lambda_1}\right) \le \mu_2/\mu_1$.
$\mathcal{P}$: $\{(\lambda_1, \lambda_2) : \lambda_2 \ge \lambda_1\} - \mathcal{I}_H$.

- $\mathcal{P}_1$: $\lambda_1 \le 1$ and $\lambda_2 \ge 1$.
- $\mathcal{P}_2$: $\lambda_1 \le \lambda_2 \le 1$.
- $\mathcal{P}_3$: $\{(\lambda_1, \lambda_2) : \lambda_2 \ge \lambda_1 \ge 1\} - \mathcal{I}_H$.

$\mathcal{I}_H$: $\lambda_2 \ge \lambda_1 \ge 1$ and $1 \le \left(1-\frac{1}{\lambda_2}\right)/\left(1-\frac{1}{\lambda_1}\right) \le \mu_2/\mu_1$.
$\mathcal{F}$: $\{(\lambda_1, \lambda_2) : \lambda_2 \le \lambda_1\} - \mathcal{I}_L$.

- $\mathcal{F}_1$: $\lambda_1 \ge 1$ and $\lambda_2 \le 1$.
- $\mathcal{F}_2$: $\{(\lambda_1, \lambda_2) : \lambda_2 \le \lambda_1 \le 1\} - \mathcal{I}_L$.
- $\mathcal{F}_3$: $\lambda_1 \ge \lambda_2 \ge 1$.

Lemma 1. *For a risk-group i, with positive iso-elastic demand elasticity λ_i, m_i satisfies:*

$$\lambda_i \gtreqless 1 \Leftrightarrow m_i \gtreqless 0.$$

Proof. This follows directly from the expression of m_i given in Equation (4.6).[16] □

Lemma 2. *For the highest risk-group n, with positive iso-elastic demand elasticity λ_n, m_n is bounded above at 1, i.e. $m_n \le 1$.*

Proof. By the political constraint, $\pi_n \le \mu_n$, and as $\lambda_n > 0$, $m_n \le 1$. □

Next note that, for the case of two risk-groups:

$$m_2 - m_1 = \frac{\pi_2}{\mu_2}\left(1-\frac{1}{\lambda_2}\right) - \frac{\pi_1}{\mu_1}\left(1-\frac{1}{\lambda_1}\right) = \left(1-\frac{1}{\lambda_1}\right)\frac{\pi_2}{\mu_1}\left[\left(\frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}}\right)\left(\frac{\mu_1}{\mu_2}\right) - \frac{\pi_1}{\pi_2}\right]. \tag{B.1}$$

And, specifically, for both $\mathcal{I}_L$ and $\mathcal{I}_H$,

$$1 \le \frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}} \le \frac{\mu_2}{\mu_1} \Rightarrow \frac{\mu_1}{\mu_2} \le \left(\frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}}\right)\left(\frac{\mu_1}{\mu_2}\right) \le 1. \tag{B.2}$$

Also, note that for any equilibrium premium regime $\underline{\pi} = (\pi_1, \pi_2)$, as π_1 increases from $\pi_1 = \mu_1$ (full risk-classification) to $\pi_1 = \pi_2$ (pooling), the ratio π_1/π_2 goes from μ_1/μ_2 to 1. So, for both $\mathcal{I}_L$ and $\mathcal{I}_H$, by intermediate value theorem, there exists a premium regime $\underline{\pi}^\star = (\pi_1^\star, \pi_2^\star)$, such that:

$$\frac{\pi_1^\star}{\pi_2^\star} = \left(\frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}}\right)\left(\frac{\mu_1}{\mu_2}\right), \tag{B.3}$$

so that Equation (B.1) becomes:

$$m_2 - m_1 = \left(1-\frac{1}{\lambda_1}\right)\frac{\pi_2}{\mu_1}\left[\frac{\pi_1^\star}{\pi_2^\star} - \frac{\pi_1}{\pi_2}\right]. \tag{B.4}$$

[16] We use the notation $\gtreqless$ in the following sense: $A \gtreqless B \Rightarrow C \gtreqless D$ is shorthand for $A > B \Rightarrow C > D$ and $A = B \Rightarrow C = D$ and $A < B \Rightarrow C < D$. A similar interpretation applies for the notation $\lesseqgtr$.

Lemma 3. *Suppose there are 2 risk-groups with risks $\mu_1 < \mu_2$ and iso-elastic demand elasticities λ_1 and λ_2 respectively. For equilibrium premium regime $\underline{\pi} = (\pi_1, \pi_2)$:*

3.1. $\mathcal{P}$: $m_2 - m_1 \geq 0$.

3.2. $\mathcal{I}_L$: $m_2 - m_1 \lesseqqgtr 0 \Leftrightarrow \frac{\pi_1}{\pi_2} \lesseqqgtr \frac{\pi_1^\star}{\pi_2^\star}$.

3.3. $\mathcal{F}$: $m_2 - m_1 \leq 0$.

3.4. $\mathcal{I}_H$: $m_2 - m_1 \gtreqqless 0 \Leftrightarrow \frac{\pi_1}{\pi_2} \lesseqqgtr \frac{\pi_1^\star}{\pi_2^\star}$.

Proof of Lemma 3.1.

$$\text{For } \mathcal{P}_1: \lambda_1 \leq 1 \text{ and } \lambda_2 \geq 1 \Rightarrow m_1 \leq 0 \text{ and } 0 < m_2 \leq 1 \Rightarrow m_2 - m_1 \geq 0.$$

$$\text{For } \mathcal{P}_2: \lambda_1 \leq \lambda_2 \leq 1 \Rightarrow 1 - \frac{1}{\lambda_1} \leq 1 - \frac{1}{\lambda_2} \leq 0 \text{ and } \frac{\pi_1}{\mu_1} \geq \frac{\pi_2}{\mu_2} \geq 0 \Rightarrow m_2 - m_1 \geq 0.$$

$$\text{For } \mathcal{P}_3: 1 \leq \lambda_1 \leq \lambda_2 \text{ and } \frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}} \geq \frac{\mu_2}{\mu_1} \Rightarrow \left(\frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}}\right)\left(\frac{\mu_1}{\mu_2}\right) \geq 1,\ \frac{\pi_1}{\pi_2} \leq 1 \text{ and } \lambda_1 \geq 1$$

$$\Rightarrow m_2 - m_1 = \left(1 - \frac{1}{\lambda_1}\right)\frac{\pi_2}{\mu_1}\left[\left(\frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}}\right)\left(\frac{\mu_1}{\mu_2}\right) - \frac{\pi_1}{\pi_2}\right] \geq 0.$$

So, for $\mathcal{P}$: $m_2 - m_1 \geq 0$. □

Proof of Lemma 3.2. For $\mathcal{I}_L$, $\lambda_1 \leq 1$, so:

$$m_2 - m_1 = \left(1 - \frac{1}{\lambda_1}\right)\frac{\pi_2}{\mu_1}\left[\frac{\pi_1^\star}{\pi_2^\star} - \frac{\pi_1}{\pi_2}\right] \lesseqqgtr 0 \Leftrightarrow \frac{\pi_1}{\pi_2} \lesseqqgtr \frac{\pi_1^\star}{\pi_2^\star} \quad \square$$

Proof of Lemma 3.3.

$$\text{For } \mathcal{F}_1: \lambda_1 \geq 1 \text{ and } \lambda_2 \leq 1 \Rightarrow m_1 \geq 0 \text{ and } m_2 \leq 0 \Rightarrow m_2 - m_1 \leq 0.$$

$$\text{For } \mathcal{F}_3: \lambda_1 \geq \lambda_2 \geq 1 \Rightarrow 0 \leq 1 - \frac{1}{\lambda_2} \leq 1 - \frac{1}{\lambda_1} \text{ and } \frac{\pi_2}{\mu_2} \leq \frac{\pi_1}{\mu_1} \Rightarrow m_2 - m_1 \leq 0.$$

$$\text{For } \mathcal{F}_2: \lambda_2 \leq \lambda_1 \leq 1 \text{ and } \frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}} \geq \frac{\mu_2}{\mu_1} \Rightarrow \left(\frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}}\right)\left(\frac{\mu_1}{\mu_2}\right) \geq 1,\ \frac{\pi_1}{\pi_2} \leq 1 \text{ and } \lambda_1 \leq 1$$

$$\Rightarrow m_2 - m_1 = \left(1 - \frac{1}{\lambda_1}\right)\frac{\pi_2}{\mu_1}\left[\left(\frac{1-\frac{1}{\lambda_2}}{1-\frac{1}{\lambda_1}}\right)\left(\frac{\mu_1}{\mu_2}\right) - \frac{\pi_1}{\pi_2}\right] \leq 0.$$

So, for $\mathcal{F}$: $m_2 - m_1 \leq 0$. □

Proof of Lemma 3.4. For $\mathcal{I}_H$, $\lambda_1 \geq 1$, so:

$$m_2 - m_1 = \left(1 - \frac{1}{\lambda_1}\right)\frac{\pi_2}{\mu_1}\left[\frac{\pi_1^\star}{\pi_2^\star} - \frac{\pi_1}{\pi_2}\right] \gtreqqless 0 \Leftrightarrow \frac{\pi_1}{\pi_2} \lesseqqgtr \frac{\pi_1^\star}{\pi_2^\star} \quad \square$$

Theorem 2. *Suppose there are two risk-groups with risks $\mu_1 < \mu_2$ and iso-elastic demand elasticities λ_1 and λ_2 respectively.*

2.1. $\mathcal{P}$: *Loss coverage is maximised by pooling and minimised by full risk-classification, while partial risk-classification is intermediate.*

2.2. $\mathcal{I}_L$: *Loss coverage is maximised by either pooling or full risk-classification depending on the population proportions and fair-premium demands.*

2.3. $\mathcal{F}$: *Loss coverage is maximised by full risk-classification regime and minimised by pooling, while partial risk-classification is intermediate.*

2.4. $\mathcal{I}_H$: *Loss coverage is maximised by a specific partial risk-classification regime.*

For the proof of Theorem 2, recall from Equation (4.4):

$$\Delta C = T\,\left(m_2 - m_1\right)\,\Delta\pi_1,\ \text{where}\ T > 0, \tag{B.5}$$

and the behaviour of $(m_2 - m_1)$ is outlined in Lemma 3.

Proof of Theorem 2.1. For $\mathcal{P}$: $m_2 - m_1 \geq 0$. So, as π_1 increases from full risk-classification ($\pi_1 = \mu_1$) to pooling ($\pi_1 = \pi_2$), the loss coverage, C, is an increasing function. Hence, loss coverage is maximum for pooled equilibrium and minimum for full risk-classification. Partial risk-classification is intermediate. □

Proof of Theorem 2.2. For $\mathcal{I}_L$:

$$\frac{\Delta C}{\Delta\pi_1} = T\,\left(m_2 - m_1\right) \lesseqqgtr 0 \Leftrightarrow \frac{\pi_1}{\pi_2} \lesseqqgtr \frac{\pi_1^\star}{\pi_2^\star}. \tag{B.6}$$

This implies that as π_1 increases from full risk-classification ($\pi_1 = \mu_1$) to pooling ($\pi_1 = \pi_2$), the loss coverage, C, first decreases and then increases reaching a minimum at $\pi_1 = \pi_1^\star$. Hence, loss coverage is maximum at either of the two extremes, pooled or full risk-classification. □

Proof of Theorem 2.3. For $\mathcal{F}$: $m_2 - m_1 \leq 0$. So, as π_1 increases from full risk-classification ($\pi_1 = \mu_1$) to pooling ($\pi_1 = \pi_2$), the loss coverage, C, is a decreasing function. Hence, loss coverage is maximum for full risk-classification and minimum for pooled equilibrium. Partial risk-classification is intermediate. □

Proof of Theorem 2.4. For $\mathcal{I}_H$:

$$\frac{\Delta C}{\Delta\pi_1} = T\,\left(m_2 - m_1\right) \gtreqqless 0 \Leftrightarrow \frac{\pi_1}{\pi_2} \lesseqqgtr \frac{\pi_1^\star}{\pi_2^\star}. \tag{B.7}$$

This implies that as π_1 increases from full risk-classification ($\pi_1 = \mu_1$) to pooling ($\pi_1 = \pi_2$), the loss coverage, C, first increases and then decreases reaching a maximum at $\pi_1 = \pi_1^\star$. Hence, loss coverage is maximum at the partial risk-classification regime $\underline{\pi}^\star$. □

Corollary 2.1. *In $\mathcal{I}_L$, the curve (λ_1, λ_2) producing exactly the same loss coverage for pooling and full risk-classification is given by the following parametric equation:*

$$\lambda_1 = \phi_1(\pi_0) = \frac{\log\left(\frac{a_1(\mu_2-\pi_0)}{\pi_0}\right)}{\log\left(\frac{\mu_1}{\pi_0}\right)}\ \textit{where}\ a_1 = \frac{p_1\,\tau_1\,\mu_1 + p_2\,\tau_2\,\mu_2}{p_1\,\tau_1\,(\mu_2-\mu_1)} \tag{B.8}$$

$$\lambda_2 = \phi_2(\pi_0) = \frac{\log\left(\frac{a_2(\pi_0-\mu_1)}{\pi_0}\right)}{\log\left(\frac{\mu_2}{\pi_0}\right)}\ \textit{where}\ a_2 = \frac{p_1\,\tau_1\,\mu_1 + p_2\,\tau_2\,\mu_2}{p_2\,\tau_2\,(\mu_2-\mu_1)} \tag{B.9}$$

Proof. For two risk-groups, the equilibrium condition under pooling, requires that:

$$p_1\,\tau_1\left(\frac{\mu_1}{\pi_0}\right)^{\lambda_1}(\pi_0-\mu_1) + p_2\,\tau_2\left(\frac{\mu_2}{\pi_0}\right)^{\lambda_2}(\pi_0-\mu_2) = 0. \tag{B.10}$$

If loss coverage under pooling and full risk-classification also need to be equal, then:

$$p_1\,\tau_1\left(\frac{\mu_1}{\pi_0}\right)^{\lambda_1}\mu_1 + p_2\,\tau_2\left(\frac{\mu_2}{\pi_0}\right)^{\lambda_2}\mu_2 = p_1\,\tau_1\,\mu_1 + p_2\,\tau_2\,\mu_2. \tag{B.11}$$

From Equations (B.10) and (B.11), eliminating $\left(\frac{\mu_1}{\pi_0}\right)^{\lambda_1}$ and $\left(\frac{\mu_2}{\pi_0}\right)^{\lambda_2}$, one at a time, gives:

$$\left(\frac{\mu_2}{\pi_0}\right)^{\lambda_2}\frac{\pi_0}{\pi_0-\mu_1} = a \Rightarrow \lambda_2 = \frac{\log\left(\frac{a_2(\pi_0-\mu_1)}{\pi_0}\right)}{\log\left(\frac{\mu_2}{\pi_0}\right)},\ \text{where}\ a_2 = \frac{p_1\,\tau_1\,\mu_1 + p_2\,\tau_2\,\mu_2}{p_2\,\tau_2\,(\mu_2-\mu_1)}; \tag{B.12}$$

$$\left(\frac{\mu_1}{\pi_0}\right)^{\lambda_1}\frac{\pi_0}{\mu_2-\pi_0} = b \Rightarrow \lambda_1 = \frac{\log\left(\frac{a_1(\mu_2-\pi_0)}{\pi_0}\right)}{\log\left(\frac{\mu_1}{\pi_0}\right)},\ \text{where}\ a_1 = \frac{p_1\,\tau_1\,\mu_1 + p_2\,\tau_2\,\mu_2}{p_1\,\tau_1\,(\mu_2-\mu_1)}. \tag{B.13}$$

So λ_2 is related to λ_1 through the pooled premium π_0, and can be presented as a parametric equation: $(\lambda_1, \lambda_2) = \left(\phi_1(\pi_0), \phi_2(\pi_0)\right)$. And, as $\lambda_1 < 1$ in $\mathcal{I}_L$:

$$\frac{d\lambda_2}{d\pi_0} = \frac{\mu_1 + \lambda_2\,(\pi_0-\mu_1)}{\pi_0\,(\pi_0-\mu_1)\log\left(\frac{\mu_2}{\pi_0}\right)} > 0\ \text{and}\ \frac{d\lambda_1}{d\pi_0} = \frac{\mu_2\,(1-\lambda_1) + \lambda_1\,\pi_0}{\pi_0(\mu_2-\pi_0)\log\left(\frac{\pi_0}{\mu_1}\right)} > 0. \tag{B.14}$$

In other words, in $\mathcal{I}_L$, the (λ_1, λ_2) curve demarcating the regions between optimal pooling and full risk-classification regimes, is a well-defined increasing function, as $\frac{d\lambda_2}{d\lambda_1} > 0$. □

Appendix C. Theorem 3

Theorem 3. *Suppose there are n risk-groups, with risks $\mu_1 < \mu_2 < \cdots < \mu_n$ and iso-elastic demand elasticities $\lambda_1, \lambda_2, \dots, \lambda_n$ respectively.*

Under a price collar κ, let $\underline{\pi}$ be the Nash equilibrium premium regime, sub-dividing the risk-groups into three super-groups $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{H}$, where all risk-groups in $\mathcal{L}$ pay the same premium π_L, all risk-groups in $\mathcal{H}$ pay the same premium π_H, and all risk-groups in $\mathcal{M}$ pay their fair actuarial premiums. Let μ_L and μ_H be the pooled equilibrium premiums of the risk-groups in $\mathcal{L}$ and $\mathcal{H}$ respectively. Further suppose:

$$\lambda_i = \begin{cases} \lambda_L & \text{if } i \in \mathcal{L}; \\ \lambda_H & \text{if } i \in \mathcal{H}. \end{cases}$$

Then the sensitivity of the loss coverage is given by:

$$\Delta C = T\,\left(m_H - m_L\right)\,\Delta\pi_L, \tag{5.1}$$

$$\text{where } T = \frac{p_L\,\tau_L\,\lambda_L\left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L+1}}{1-m_H}; \text{ and } m_i = \frac{\pi_i}{\mu_i}\left(1-\frac{1}{\lambda_i}\right); \quad i = L, H; \tag{5.2}$$

where for the super-groups $\mathcal{L}$ and $\mathcal{H}$, p_L and p_H are the aggregate proportion of population belonging to the super-groups; and τ_L and τ_H can be interpreted as the 'fair-premium demand' when all risk-groups in the respective super-groups are pooled and charged the same pooled premium.

Proof. As the risk-groups in $\mathcal{M}$ do not contribute to profit or loss, the equilibrium condition can be expressed as:

$$E\left(\underline{\pi}\right) = \underbrace{\sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_i}{\pi_L}\right)^{\lambda_L}\left(\pi_L - \mu_i\right)}_{E_{\mathcal{L}}} + \underbrace{\sum_{j\in\mathcal{H}} p_j\,\tau_j\left(\frac{\mu_j}{\pi_H}\right)^{\lambda_H}\left(\pi_H - \mu_j\right)}_{E_{\mathcal{H}}} = 0. \tag{C.1}$$

The first term, $E_{\mathcal{L}}$, in Equation (C.1), can be split as follows:

$$E_{\mathcal{L}} = \sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_i}{\pi_L}\right)^{\lambda_L}\left(\pi_L - \mu_i\right); \tag{C.2}$$

$$= \sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L}\left(\frac{\mu_i}{\mu_L}\right)^{\lambda_L}\left[\left(\pi_L - \mu_L\right) + \left(\mu_L - \mu_i\right)\right]; \tag{C.3}$$

$$= \left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L}\left[\sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_i}{\mu_L}\right)^{\lambda_L}\left(\pi_L - \mu_L\right) + \sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_i}{\mu_L}\right)^{\lambda_L}\left(\mu_L - \mu_i\right)\right]; \tag{C.4}$$

where μ_L is such that the second term in Equation (C.4), is zero, i.e.:

$$\sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_i}{\mu_L}\right)^{\lambda_L}\left(\mu_L - \mu_i\right) = 0, \tag{C.5}$$

so that μ_L can be interpreted as the pooled equilibrium premium, if the insurance market only consisted of the risk-groups in $\mathcal{L}$. Also, μ_L is unique and is given by:

$$\mu_L = \frac{\sum_{i\in\mathcal{L}} p_i\,\tau_i\,\mu_i^{\lambda_L+1}}{\sum_{i\in\mathcal{L}} p_i\,\tau_i\,\mu_i^{\lambda_L}}, \quad \text{so that:} \quad \mu_1 \le \mu_L \le \max_{i\in\mathcal{L}} \mu_i \le \pi_L. \tag{C.6}$$

Using such a μ_L, the expression for $E_{\mathcal{L}}$ in Equation (C.4) becomes:

$$E_{\mathcal{L}} = \left[\sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_i}{\mu_L}\right)^{\lambda_L}\right]\left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L}\left(\pi_L - \mu_L\right) = p_L\tau_L\left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L}\left(\pi_L - \mu_L\right); \tag{C.7}$$

$$\text{where } p_L = \sum_{i\in\mathcal{L}} p_i,\ \tau_L = \sum_{i\in\mathcal{L}}\left(\frac{p_i}{p_L}\right)\tau_i\left(\frac{\mu_i}{\mu_L}\right)^{\lambda_L}. \tag{C.8}$$

Note that p_L is the aggregate proportion of population belonging to the collection of risk-groups in $\mathcal{L}$ and τ_L can be interpreted as the 'fair-premium demand' when all risk-groups in $\mathcal{L}$ are pooled and charged the same pooled premium, μ_L.

A similar line of argument for the risk-groups in $\mathcal{H}$ leads to:

$$E_{\mathcal{H}} = p_H\tau_H\left(\frac{\mu_H}{\pi_H}\right)^{\lambda_H}\left(\pi_H - \mu_H\right); \text{ where } p_H = \sum_{j\in\mathcal{H}} p_j,\ \tau_H = \sum_{j\in\mathcal{H}}\left(\frac{p_j}{p_H}\right)\tau_j\left(\frac{\mu_j}{\mu_H}\right)^{\lambda_H}, \tag{C.9}$$

$$\text{where } \mu_H = \frac{\sum_{j\in\mathcal{H}} p_j\,\tau_j\,\mu_j^{\lambda_H+1}}{\sum_{j\in\mathcal{H}} p_j\,\tau_j\,\mu_j^{\lambda_H}}, \quad \text{so that:} \quad \pi_H \le \min_{j\in\mathcal{H}} \mu_j \le \mu_H \le \mu_n. \tag{C.10}$$

Using the expressions for $E_{\mathcal{L}}$ and $E_{\mathcal{H}}$ in Equations (C.7) and (C.9) respectively, Equation (C.1) becomes:

$$E\left(\underline{\pi}\right) = \underbrace{p_L\,\tau_L\left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L}\left(\pi_L - \mu_L\right)}_{E_{\mathcal{L}}} + \underbrace{p_H\,\tau_H\left(\frac{\mu_H}{\pi_H}\right)^{\lambda_H}\left(\pi_H - \mu_H\right)}_{E_{\mathcal{H}}} = 0. \tag{C.11}$$

Equation (C.11) shows that it is possible to conceptualise $\mathcal{L}$ and $\mathcal{H}$ as super-groups with demand elasticities λ_L and λ_H respectively, where the true risks of the super-groups are taken to be pooled equilibrium premiums, μ_L and μ_H, of the respective super-groups. Essentially, this reduces the problem involving more than two risk-groups to the simpler two risk-groups problem, so that the analysis of Section 4 can be extended directly to this situation.

The loss coverage for the Nash premium regime can then be expressed as:

$$C\left(\underline{\pi}\right) = \sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_i}{\pi_L}\right)^{\lambda_L}\mu_i + \sum_{j\in\mathcal{H}} p_j\,\tau_j\left(\frac{\mu_j}{\pi_H}\right)^{\lambda_H}\mu_j + \sum_{m\in\mathcal{M}} p_m\,\tau_m\,\mu_m; \tag{C.12}$$

$$= \sum_{i\in\mathcal{L}} p_i\,\tau_i\left(\frac{\mu_i}{\pi_L}\right)^{\lambda_L}\pi_L + \sum_{j\in\mathcal{H}} p_j\,\tau_j\left(\frac{\mu_j}{\pi_H}\right)^{\lambda_H}\pi_H + \sum_{m\in\mathcal{M}} p_m\,\tau_m\,\mu_m; \tag{C.13}$$

... by the equilibrium condition in Equation (C.1);

$$= p_L\,\tau_L\left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L}\pi_L + p_H\,\tau_H\left(\frac{\mu_H}{\pi_H}\right)^{\lambda_H}\pi_H + \sum_{m\in\mathcal{M}} p_m\,\tau_m\,\mu_m; \tag{C.14}$$

... by the definitions of p_L, τ_L, μ_L, p_H, τ_H and μ_H;

$$= p_L\,\tau_L\left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L}\mu_L + p_H\,\tau_H\left(\frac{\mu_H}{\pi_H}\right)^{\lambda_H}\mu_H + \sum_{m\in\mathcal{M}} p_m\,\tau_m\,\mu_m; \tag{C.15}$$

by the equilibrium condition in Equation (C.11).

Assuming the compositions of $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{H}$ remain unaffected, changing π_L affects π_H without any implications for the risk-groups in $\mathcal{M}$. Also note that, as long as $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{H}$ remain unchanged, we can follow the same steps as in Section 4, to get:

$$\Delta C = T\,\left(m_H - m_L\right)\Delta\pi_L, \tag{5.1}$$

$$\text{where } T = \frac{p_L\,\tau_L\,\lambda_L\left(\frac{\mu_L}{\pi_L}\right)^{\lambda_L+1}}{1-m_H}; \text{ and } m_i = \frac{\pi_i}{\mu_i}\left(1-\frac{1}{\lambda_i}\right); \quad i = L, H. \quad \square \tag{5.2}$$

References

Akerlof, G.A., 1970. The market for lemons: quality uncertainty and the market mechanism. The Quarterly Journal of Economics 84, 488–500.

Auerbach, A.J., Kotlikoff, L.J., 1991. The adequacy of life insurance purchases. Journal of Financial Intermediation 1 (3), 215–241.

Babbel, D., 1985. The price elasticity of demand for whole life insurance. The Journal of Finance 40 (1), 225–239.

Barseghyan, L., Molinari, F., O'Donoghue, T., Teitelbaum, J.C., 2013. The nature of risk preferences: evidence from insurance choices. American Economic Review 103 (6), 2499–2529.

Bernheim, B.D., Formi, L., Gokhale, J., Kotlikoff, L.J., 2003. The mismatch between life insurance holdings and financial vulnerabilities: evidence from the health and retirement study. American Economic Review 93 (1), 354–365.

Bhargava, S., Loewenstein, G., Sydnor, J., 2017. Choose to lose: health plan choices from a menu with dominated options. The Quarterly Journal of Economics 132 (3), 1319–1372.

Blumberg, L., Nichold, L., Banthin, J., 2001. Worker decisions to purchase health insurance. International Journal of Health Care Finance and Economics 1, 305–325.

Buchmueller, T.C., Ohri, S., 2006. Health insurance take-up by the near-elderly. Health Services Research 41, 2054–2073.

Butler, J.R., 1999. Estimating elasticities of demand for private health insurance in Australia. Working Paper No. 43. National Centre for Epidemiology and Population Health, Australian National University, Canberra.

Chatterjee, I., Macdonald, A.S., Tapadar, P., Thomas, R.G., 2021. When is utilitarian welfare higher under insurance risk pooling? Insurance: Mathematics and Economics (ISSN 0167-6687). https://doi.org/10.1016/j.insmatheco.2021.08.006. https://www.sciencedirect.com/science/article/pii/S0167668721001268.

Chernew, M., Frick, K., McLaughlin, C., 1997. The demand for health insurance coverage by low-income workers: Can reduced premiums achieve full coverage? Health Services Research 32, 453–470.

Dixit, A.K., 1990. Optimization in Economic Theory. Oxford University Press. ISBN 9780198772101.

Dosis, A., 2022. Price caps and efficiency in markets with adverse selection. Journal of Mathematical Economics 99, 102591. https://doi.org/10.1016/j.jmateco.2021.102591.

Einav, L., Finkelstein, A., 2023. Empirical analyses of selection and welfare in insurance markets: a self-indulgent survey. The Geneva Risk and Insurance Review 48 (2), 167–191. https://doi.org/10.1057/s10713-023-00084-3.

Einav, L., Finkelstein, A., Cullen, M.R., 2010. Estimating welfare in insurance markets using variation in prices. The Quarterly Journal of Economics 125 (3), 877–921.

Einav, L., Finkelstein, A., Tebaldi, P., 2019. Market design in regulated health insurance markets: Risk adjustment vs. subsidies. Working Paper.

Ericson, K.M.M., 2014. Consumer inertia and firm pricing in the Medicare Part D prescription drug insurance exchange. American Economic Journal: Economic Policy 6 (1), 38–64.

Ericson, K.M.M., Starc, A., 2015. Pricing regulation and imperfect competition on the Massachusetts health insurance exchange. Review of Economics and Statistics 97 (3), 667–682.

Geruso, M., Layton, T.J., McCormack, G., Shepard, M., 2023. The two margin problem in insurance markets. Review of Economics and Statistics 105 (2), 237–257.

Goda, G.S., 2011. The impact of state tax subsidies for private long-term care insurance on coverage and medicaid expenditures. Journal of Public Economics 95, 744–757.

Goodwin, B.K., 1993. An empirical analysis of the demand for multiple peril crop insurance. American Journal of Agricultural Economics 75 (2), 424–434.

Gruber, J., Sommers, B.D., 2019. The Affordable Care Act's effects on patients, providers and the economy; what we've learned so far. Journal of Policy Analysis and Management 38 (4), 1028–1052.

Handel, B.R., 2013. Adverse selection and inertia in health insurance markets: when nudging hurts. American Economic Review 103 (7), 2643–2682.

Handel, B.R., Kolstad, J.T., 2015. Health insurance for humans: information frictions, plan choice, and consumer welfare. American Economic Review 105 (8), 2449–2500.

Hao, M., Macdonald, A.S., Tapadar, P., Thomas, R.G., 2016. Insurance loss coverage under restricted risk classification: The case of iso-elastic demand. ASTIN Bulletin 46 (2), 265–291. https://doi.org/10.1017/asb.2016.6.

Hao, M., Macdonald, A.S., Tapadar, P., Thomas, R.G., 2018. Insurance loss coverage and demand elasticities. Insurance: Mathematics and Economics 79, 15–25. https://doi.org/10.1016/j.insmatheco.2017.12.002.

Hao, M., Macdonald, A.S., Tapadar, P., Thomas, R.G., 2019. Insurance loss coverage and social welfare. Scandinavian Actuarial Journal 2019 (2), 113–128. https://doi.org/10.1080/03461238.2018.1513865.

Hendren, N., 2021. Measuring ex-ante welfare in insurance markets. The Review of Economic Studies 88 (3), 1193–1223.

Hoy, M., 2006. Risk classification and social welfare. The Geneva Papers on Risk and Insurance. Issues and Practice 31, 245–269.

Lindholm, M., Richman, R., Tsanakas, A., Wüthrich, M.V., 2021. Discrimination-free insurance pricing. ASTIN Bulletin 52 (2), 55–89.

Mahoney, N., Weyl, E.G., 2017. Imperfect competition in selection markets. Review of Economics and Statistics 99 (4), 637–651.

Malpas, P.J., 2008. Is genetic information relevantly different from other kinds of non-genetic information in the life insurance context? Journal of Medical Ethics 34 (7), 548–551.

Mittra, J., 2006. Genetic exceptionalism and precautionary politics: regulating for uncertainty in Britain's genetics and insurance policy process. Science and Public Policy 33 (8), 585–600.

Murray, T.H., 2019. Is genetic exceptionalism past its sell-by date? The American Journal of Bioethics 19 (1), 13–15.

Pauly, M.V., Withers, K.H., Viswanathan, K.S., Lemaire, J., Hershey, J.C., Armstrong, K., Asch, D.A., 2003. Price elasticity of demand for term life insurance and adverse selection. NBER Working Paper (9925).

Pope, D.G., Sydnor, J.R., 2011. Implementing anti-discrimination policies in statistical profiling models. American Economic Journal: Economic Policy 3 (3), 206–231.

Rothschild, M., Stiglitz, J., 1976. Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. The Quarterly Journal of Economics 90 (4), 630–649.

Tebaldi, P., 2022. Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the ACA. Working Paper.

Thomas, R.G., 2008. Loss coverage as a public policy objective for risk classification schemes. The Journal of Risk and Insurance 75, 997–1018.

Thomas, R.G., 2009. Demand elasticity, risk classification and loss coverage: when can community rating work? ASTIN Bulletin 39, 403–428.

Thomas, R.G., 2017. Loss Coverage: Why Insurance Works Better with Some Adverse Selection. Cambridge University Press.

Vazquez, A., 1995. A note on arc elasticity of demand. Estudios Economicos 10 (2), 221–228.

Viswanathan, K.S., Lemaire, J., Withers, K.K., Armstrong, K., Baumritter, A., Hershey, J., Pauly, M., Asch, D.A., 2006. Adverse selection in term life insurance purchasing due to the BRCA 1/2 genetic test and elastic demand. The Journal of Risk and Insurance 74, 65–86.

Wettstein, G., 2017. The Affordable Care Act's insurance market regulations effect on coverage. Health Economics 27 (3), 454–464.

Wilson, C., 1977. A model for insurance with incomplete information. Journal of Economic Theory 16, 167–207.