ISSN 0956-8549

Market Information and the Impact of Public Health Insurance on the Private Insurance Market

By

Alistair McGuire

David Webb

FINANCIAL MARKETS GROUP DISCUSSION PAPER NO. 964

July 2026

Market Information and the Impact of Public Health Insurance on the Private Insurance Market

Alistair McGuire and David Webb
London School of Economics

April 17, 2026

Alistair McGuire, (corresponding author),
Professor of Health Economics,
Department of Health Policy,
London School of Economics and Political Science,
Houghton Street, London WC2A 2AE, United Kingdom
Email: a.j.mcguire@lse.ac.uk
and
David Webb, Professor of Finance,
Department of Finance,
London School of Economics and Political Science,
Houghton Street, London WC2A 2AE, United Kingdom Email: d.c.webb@lse.ac.uk

There are no conflicts of interest or declarations of financial support.

Each author was involved in conceptualisation of the project, the formal analysis, developing the methodology, writing the original draft and reviewing and editing

Acknowledgements. We thank the attendees of the LSE London Health Economics group for their comments on an earlier version and Marcos Vera-Hernández, UCL.

Abstract

We analyse how public health insurance affects private insurance markets in mixed public-private systems. Extending the standard selection framework to incorporate heterogeneous risk aversion, quality preferences, income constraints, and administrative costs, we characterise equilibrium outcomes under both adverse and advantageous selection, and identify a double-margin mechanism that arises under the latter. Under adverse selection, public coverage attracts low-risk types away from the private market, worsening the private risk pool, raising premiums, and generating further exit; a self-reinforcing deterioration of private contract terms operating solely at the intensive margin. Under advantageous selection the direction reverses and public coverage now attracts high-risk, low quality-preference types, endogenously improving private contract terms through the zero-profit condition. Under institutional structures where public insurance is means-tested and private insurance is priced above actuarial value due to administrative costs, this resultant improvement in the private insurance pool crowds-in previously excluded marginal types at the extensive margin. We show that the two margins, (intensive and extensive) are causally linked. This feedback loop is absent from existing models. For the extensive margin effect the mechanism requires exactly three conditions to hold: advantageous selection, administrative costs, and income-based eligibility for public insurance. Institutional setting is shown to be important. The double-margin operates fully under means-tested public insurance, and takes a modified form under mandatory basic insurance with voluntary supplementary coverage. Only the intensive margin operates under NHS-type universal coverage. Standard crowd-out measures, which capture only the intensive margin under either selection regime systematically misstate the consequences of public insurance expansion.

1 Introduction

Public and private health insurance coexist in most advanced economies, yet the interactions between the two sectors remain only partially understood. The dominant view is that expanding public coverage crowds out private insurance, as individuals shift to public coverage and exit the private market.[1]

This paper identifies a fuller mechanism, which we refer to as the double-margin mechanism, through which expanding public coverage simultaneously generates crowd-out of high-risk of illness individuals and crowd-in of low-risk individuals at the intensive margin (the margin operating on those individuals who already hold insurance coverage, whether through public or private provision) as the initial crowd-out improves the private pool. Under certain institutional structures further crowd-in at the extensive margin, (of marginal individuals who now buy private insurance in reaction to a change in the premium), will also occur giving rise to our double-margin (intensive and extensive) effect. The crowd-in mechanisms have been overlooked in the existing literature. These internal and external margin effects are not independent but are causally linked. It is precisely the crowd-out of high-risk types that improves the terms of private contracts and draws previously excluded individuals into private coverage. Standard crowd-out measures operating only on the intensive margin systematically miscalculate the consequences of public insurance expansion regardless of the selection regime. We also show that the direction of the resulting bias can differ, as under adverse selection crowd-out worsens private contract terms, while under advantageous selection it improves them.

We demonstrate the double-margin mechanism in a mixed system where means-tested public insurance coexists with a competitive private market - an environment broadly analogous to Medicaid in the United States. In this setting we show that the resulting equilibrium requires three jointly necessary conditions (identified formally in Proposition 1): advantageous selection, administrative costs, and income-based eligibility for public insurance. Remove any one and the mechanism collapses to a single margin: without advantageous

[1]The empirical literature consistently reports crowding-out. Cutler and Gruber (1996) and Gruber and Simon (2008) estimate that the effect is large: up to 60% of Medicaid expansion was offset by reductions in private coverage. Others find more modest effects, with most estimates in the range of 10% to 20% (Dubay and Kenney, 1996, 1997; Thorpe and Florence, 1998; Shore-Sheppard et al., 2000; Card and Shore-Sheppard, 2004). These empirical estimates capture the intensive margin only and measure the net displacement of privately insured individuals into public coverage. They do not separately identify any offsetting crowd-in at either margin. Under the double-margin mechanism, such estimates will systematically overstate the net reduction in private coverage to the extent that crowd-in effects are present but undetected. Simulation models yield mixed findings: Golosov and Tsyvinski (2015) argue that crowding-out may be substantial, while Chetty and Saez (2017) find that outcomes depend on the distribution of under- and over-insurance within a population.

selection the private premium rises rather than falls; without administrative costs no individuals are priced out of the private market by premium loading; and without income-based eligibility all risk types can access public coverage directly. We illustrate this by applying the model to a different institutional structure (characterised by the UK NHS) where public insurance provision is not means-tested and the third condition does not hold, such that the crowd-out and crowd-in effects operate only on the intensive margin.

This paper's contribution spans two strands of the literature, and its relationship to each is distinct. The first is the theory of selection in insurance markets. The classic Rothschild and Stiglitz (1976) adverse selection framework arsing from information asymmetry, supports a separating equilibrium, while pooling is shown to be unstable. De Meza and Webb (2001) establish, however, that when risk aversion and illness risk are negatively correlated reflecting advantageous selection, a stable pooling equilibrium can exist. The present paper extends these models to a mixed public-private system showing that public expansion can lead to crowd-out, but also trigger endogenous improvements in private contract terms sufficient to activate further crowd-in effects.[2] This distinction matters in its own right, as the classic literature does not consider the coexistence of a public insurance offer alongside private insurance, even though this is a near-universal feature of modern health insurance markets.

The second strand concerns the extensive and intensive margins in insurance markets. Geruso, Layton, McCormack and Shepard (2023), focusing solely on adverse selection, demonstrate that the extensive and intensive margins are interdependent, with policies operating on each affecting distinct marginal populations through endogenous price changes. Weyl and Veiga (2016), also restricting analysis to adverse selection, show that pricing arrangements generate systematic variation across the two margins in line with institutional differences.[3] Our contribution here is to show that in mixed public-private systems the two margins also interact causally, with the crowd-out of high-risk of illness types generating an endogenous improvement in private contract terms that can generate crowd-in at the intensive margin (see Sections 2.2 and 3.2), but that also activates crowd-in at the extensive margin in a means-tested public system (see Section 4.1) - a feedback mechanism that is

[2] Olivella and Vera-Hernández (2013), in an appendix, consider an endogenous feedback between private sector provision operating through taxation on income and subsequent private insurance take-up, but do not identify the crowd-in effect.

[3] This strand of the literature, with its focus on endogenous price change effects across different marginal populations, is close in spirit to the present paper. Our emphasis differs, however: we treat public and private providers as operating on the risk pool itself, rather than on endogenous price changes per se, in keeping with the original work on adverse selection and advantageous selection. We also consider a wider range of selection and institutional arrangements as they impact on the risk pools. Work on the impact of endogenous quality impacts on insurance demand is also related, though less directly so (see Chade et al, 2024; Veiga and Weyl, 2016).

absent from existing analysis.[4]

The double-margin mechanism has direct policy implications. Evaluating public insurance expansions purely in terms of individuals displaced from private coverage, as much of the empirical work on Medicaid expansion has done, omits the potential crowd-in margins from the associated welfare gains arising from extending coverage to previously uninsured types. Whether the crowd-in is large enough to offset the fiscal cost of extended public eligibility is an empirical question. The prior theoretical point, however, is whether both margins are activated and what their net welfare effect is. The institutional setting is central here as it determines the relevant policy analysis. We show that the double-margin operates fully under means-tested public insurance (analogous to Medicaid in the USA), is absent under NHS-type universal coverage as the extensive margin does not exist as all individuals are eligible for public insurance, and operates in modified form under mandatory basic insurance with voluntary supplementary coverage as in the Netherlands or the coverage patterns in the USA presented by Medicare and Medigap. These comparisons highlight that the same underlying selection forces produce qualitatively different outcomes depending on how public eligibility is structured.

The paper proceeds as follows. Section 2 introduces an innovative and intuitive graphical framework in premium-coverage space and establishes the baseline model of coexisting public and private insurance under symmetric information, adverse selection, and advantageous selection. Section 3 enriches the basic model with heterogeneous quality preferences.[5] Section 4 introduces administrative costs, risk preferences and income heterogeneity, while deriving the double-margin mechanism formally. Section 4.4 analyses the mechanism's operation across three common institutional settings, outlining a further result as it relates to institutional dependence. Section 5 concludes. Formal proofs supporting the graphical analysis and welfare analysis are collected in the Appendices.

[4] Encinosa (2003) provides a taxonomy of such effects but does not provide the formal modelling.

[5] Although this paper does not extend to empirical analysis, Sections 2 and 3 have direct methodological implications for the empirical testing of asymmetric information in insurance markets. Standard tests for asymmetric information are derived from models of either pure adverse or pure advantageous selection in a single market; in a mixed public-private system, the coexistence of the two sectors alters the composition of the private risk pool in ways that confound such tests. Olivella and Vera-Hernández (2013) make this point the focus of their analysis as applied to their empirical analysis of the UK NHS.

2 Model of the Joint Provision of Private Insurance and Public Health Cover

We model the coexistence of a public insurance sector and a competitive private insurance market with two types of individuals who differ in the probability of ill-health, θ_i. These groups are denoted as high-risk, H, and low-risk, L, with $\theta_H > \theta_L$. All individuals are risk averse, captured by a strictly concave utility function $U_i(\cdot)$ with $U_i' > 0$ and $U_i'' < 0$. The population proportion of type L is α and the proportion of type H is $(1-\alpha)$. Both groups have income Y and face the prospect of illness at total monetary cost $\bar{C}$, a fraction $0 \leq \lambda \leq 1$ of which may be covered by private insurance at premium Π. Preferences over private insurance are given by

$$V_i = \theta_i U_i\big(Y - \bar{C} + \lambda\bar{C} - \Pi_i\big) + (1-\theta_i)U_i(Y - \Pi_i), \quad i = H, L \tag{1}$$

Insurance contracts are offered by competitive insurers engaged in Bertrand competition, with premium Π_i satisfying

$$\Pi_i \geq \theta_i \lambda_i \bar{C}, \quad i = H, L \tag{2}$$

as equalities in equilibrium. The private market operates as a screening game: in stage 1, insurers make irrevocable contract offers specifying Π and payout $\lambda\bar{C}$; in stage 2, individuals choose the contract maximising expected utility given their private information. We consider only pure-strategy subgame-perfect Nash equilibria, under both symmetric and asymmetric information. In the Figures below, indifference curves are denoted I, with those for type H given by I_H and those for type L by I_L. Offer curves are denoted by O, with premium on the horizontal axis and coverage on the vertical axis., with that for the L group, O_L being steeper than that for the H group, O_H.[6] Throughout, greater convexity of an indifference curve indicates greater risk aversion. We refer to Rothschild and Stiglitz (1976), de Meza and Webb (2001), and Olivella and Vera-Hernández (2013) for further analytical details. As the equilibrium structures under both adverse and advantageous selection are well established in the literature, we invoke these standard results without reproducing the full proofs, although we do provide a concise summary of these proofs in Appendix A.

Generally we use the term crowd-out in the traditional sense to refer to an intensive margin associated with the exit of privately insured individuals into public coverage, but we also have an intensive margin associated with individuals who crowd-into the private insurance sector from the public sector as the risk profile of the population of the private market

[6]Solely for illustrative purposes the x-axis is stretched so that offer curves do not lie above a 45° line. This does not affect the underlying properties of the model.

changes. By extensive margin we refer to the entry of previously uninsured individuals into private insurance.[7] The analytical focus of the paper is the double-margin mechanism, activated through both the intensive and extensive margins, developed in Section 4, which itself builds on the advantageous selection equilibrium based on intensive margin crowd-in, that we introduce here in Section 2.2.

Under symmetric information the full-insurance contracts for H and L are at S_H and S_L respectively, located at the tangency points of each group's indifference curve with the respective offer curves at the full-loss exposure $\bar{C}$. If a non-means-tested public insurance contract with net value G is available at $(0, G)$, the H group opts for public coverage first:

$$V_H(G) = \theta_H U_H\big(Y - \bar{C} + G\big) + (1 - \theta_H)U_H(Y) \geq V_H(S_H), \quad V_L(G) < V_L(S_L) \tag{3}$$

This baseline case of coexisting public-competitive private insurance with symmetric information is illustrated in Figure 1.[8]

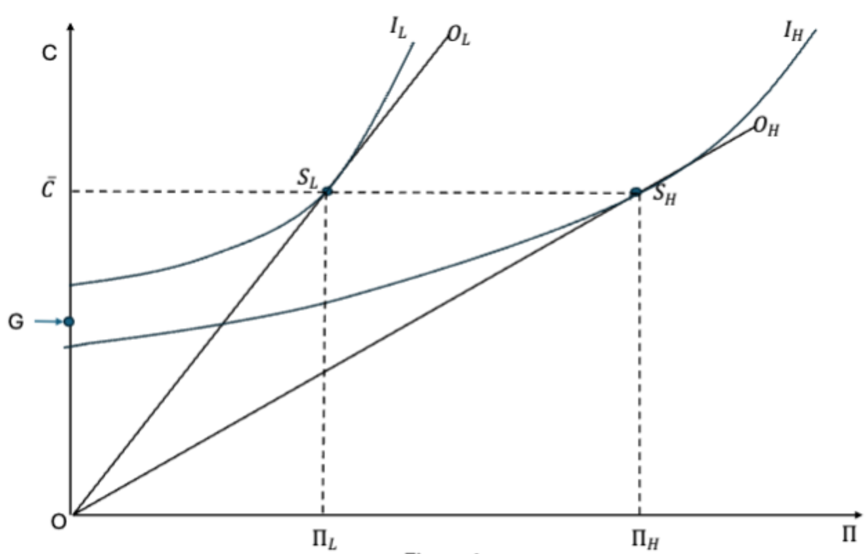

Figure 1

[7]We amend the definition of crowd-in at the intensive and extensive margins slightly when considering the Dutch insurance system below to reflect the complexity of that system

[8]Note that throughout we are not defining the optimal level of public insurance (G), but rather introducing cases of interest where the public coverage will impact on the private equilibrium.

2.1 Adverse Selection with the Coexistence of Public-Private Insurance

Now assume asymmetric information with θ_i private. With single crossing of indifference curves,

$$\frac{\partial(MRS_H)}{\partial(\lambda\bar{C})} < \frac{\partial(MRS_L)}{\partial(\lambda\bar{C})},$$

the Rothschild–Stiglitz separating equilibrium gives contracts (S_H, S'_L) satisfying

$$V_H(S_H) \geq V_H(S'_L); \qquad V_L(S'_L) \geq V_L(S_H) \tag{4}$$

with L obtaining partial coverage, $\lambda_L\bar{C}$, $\lambda_L < 1$, at S'_L; and H obtaining full coverage at S_H. The H-type incentive constraint binds; the L-type constraint is slack. The quantity $(1-\lambda_L)\bar{C}$ is a deductible. Pooling equilibria are not sustainable under these conditions: by single crossing, any pooling contract can be profitably undercut by a contract attracting only low-risk types (as shown in Appendix A.1). Conditional on the level of public coverage G, the group attracted to the public option differs from the symmetric information case. At a sufficiently generous $G_2 > G_1$,

$$V_L = \theta_L U_L\big(Y - \bar{C} + G_2\big) + (1-\theta_L)U_L(Y) \geq V_L(S'_L), \qquad V_H(G_2) < V_H(S_H) \tag{5}$$

so under asymmetric information the L group may take public coverage first, making the private pool riskier. This separating equilibrium is illustrated in Figure 2 (and formally characterised in Appendix A.1). While the equilibrium follows the classic Rothschild-Stiglitz separating structure, with the L types opting for public insurance (G_2) and the H types preferring higher private coverage as just noted the pooling equilibria (at S') are not sustainable in this single-crossing adverse selection framework. Under asymmetric information and the single-crossing property, any pooling contract is vulnerable to a profitable deviation, as insurers can always offer a contract that attracts only low-risk types and satisfies the zero-profit condition, breaking the pooling equilibrium.

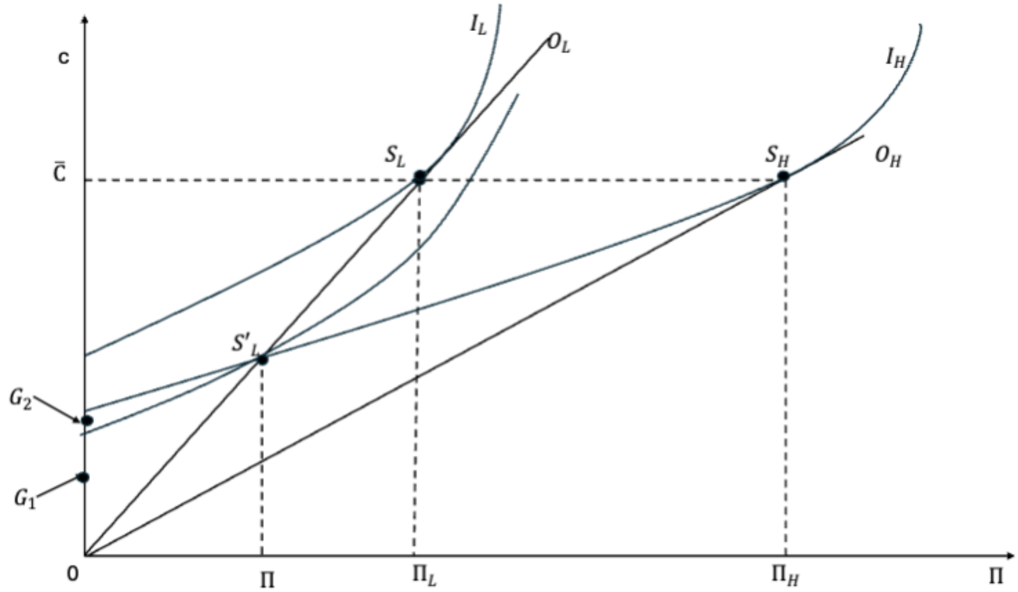

Figure 2

With symmetric information we have shown that the public health insurance option first attracts the H group. However, conversely, with asymmetric information and single crossing, the public health insurance option may, depending on the level of coverage (e.g. G_2), first attract the L group, who may prefer it over private coverage with a high deductible (in Figure 2, the difference in coverage between S_L and S'_L). So conditional on the level of public insurance, low-risk types may take-up public coverage making the private insurance pool riskier. The adverse selection case thus establishes a deterioration channel: public expansion worsens the private pool and reinforces crowd-out exit at the intensive margin. Section 2.2 now introduces advantageous selection, which reverses this direction at the intensive margin and provides the equilibrium foundation for the double-margin mechanism which incorporates an extensive margin effect in Section 4.

2.2 Advantageous Selection with the coexistence of public-private insurance

Advantageous selection challenges the prediction that higher-risk of illness individuals always purchase more generous insurance. De Meza and Webb (2001) formalise this using a negative correlation between risk aversion and illness risk, while Boone and Schottmüller (2017) provide an alternative foundation through income heterogeneity and endogenous treatment choice, showing that the single crossing condition at the heart of Rothschild and Stiglitz may be violated and that healthier, wealthier individuals seek better private coverage. The income-health correlation underlying this result is well-documented: higher-income individuals typically have better health outcomes (lower θ) and higher willingness-to-pay for coverage quality. This correlation can also arise endogenously from preventative investment choices at stage 2 of the screening game, as more risk-averse individuals accumulate greater wealth through conservative financial behaviour and maintain better health through conservative lifestyle choices, generating the wealth-health correlation needed for advantageous selection. Irrespective of the justification offered, such correlations result in a double-crossing of indifference curves. Such results provide a theoretical foundation for understanding underinsurance among high-risk individuals in the private insurance market. Empirical support is provided by Fang, Keane and Silverman (2008), who show in the context of the Medigap insurance market that advantageous selection arises from differences in risk aversion rather than risk of illness differences alone, consistent with the mechanism modelled here.

Assume asymmetric information and that L types are more risk-averse than H types. Absent a public option, the equilibrium involves fraction ε of H types taking the separating contract S_H, with the remaining fraction $(1-\varepsilon)$ pooling with all L types at contract P, where

$$\Pi_P \geq \frac{[\alpha\theta_L + (1-\alpha)(1-\varepsilon)\theta_H]}{\alpha + (1-\alpha)(1-\varepsilon)}\lambda_P \bar{C} \tag{6}$$

as an equality, and H types are indifferent between S_H and P (the private pooling equilibrium).

Figure 3 illustrates a partial pooling equilibrium with a public option with coverage at G, such that a fraction ε' of H types take public coverage and the remainder pool with all L types at P', (the private pooling equilibrium if public insurance is available a G) satisfying

$$V_H(P') = V_H(G) \tag{7}$$

The contract P' dominates P for both groups: as an increasing proportion of H types switch to public coverage the private pool becomes healthier, premiums fall and coverage improves, generating crowd-in at the intensive margin as more L types find private insurance attractive. This partial pooling equilibrium is illustrated in Figure 3 (and formally characterised in Appendix A.2).

Two features of this equilibrium are central to the paper's main result. First, stability rests on the double-crossing property: the indifference curves of H and L types cross twice, pinning down a unique pooling contract and preventing Rothschild-Stiglitz unravelling. Second, P' is endogenous to G: as G improves, ε' rises, the private pool improves, and the zero-profit condition drives contract terms upward. This endogenous improvement is the

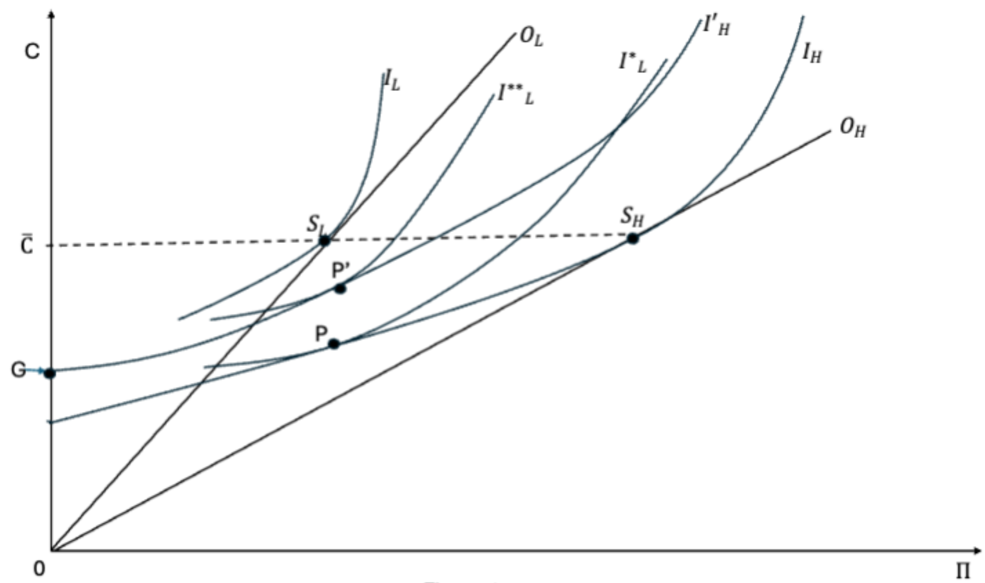

Figure 3

channel through which crowd-out at the intensive margin also generates crowd-in at the intensive margin. Sections 3 and 4 complete the understanding of this mechanism by introducing quality heterogeneity, where again with advantageous selection crowd-in operates at the intensive margin, and administrative costs and means-tested public insurance where crowd-in also occurs at the extensive margin.

3 Model with Health Risk Difference and Differences in Preference for Quality/Coverage of Health Care

We extend the framework to incorporate heterogeneity in preferences for healthcare quality.[9] Let preferences be written as

$$V_i = \theta_i U_i\big(Y - \bar{C}(q) + \lambda_i \bar{C}(q) - \Pi(q)\big) + (1-\theta_i)U_i(Y - \Pi(q)) + \theta_i \alpha_j q \tag{8}$$

where q is a quality index, $\bar{C}(q)$ is increasing in q, and $\alpha_j q$ ($j = h, l$) (with $\alpha_h > \alpha_l$) captures high and low willingness-to-pay for quality. With full information insurance would be priced fairly:

$$\Pi(q_j) \geq \theta_i \lambda_{i,j} \bar{C}(q_j), \quad i = H, L;\ j = h, l \tag{9}$$

as equalities in equilibrium.

3.1 Adverse Selection with coexisting public-private insurance and different preferences for quality

With both risk and quality heterogeneity, selection outcomes depend on the joint distribution of risk of illness type and willingness-to-pay for quality of coverage. Figure 4 illustrates the case where low risk of illness types, L, have high willingness-to-pay for quality (Lh) and high risk of illness types, H, have low willingness-to-pay for quality (Hl).[10] Under asymmetric information the Rothschild-Stiglitz separating equilibrium is at (S'_{Lh}, S_{Hl}). Once public cover is introduced at G the Lh type is attracted to this public cover first, leaving the Hl group to take full private insurance at S_{Hl}. Figure 5 shows the reverse configuration (Ll and Hh): the separating equilibrium is at (S'_{Ll}, S_{Hh}) and once public coverage is introduced at G the Ll type is attracted to public cover first. Whether high-risk or low-risk of illness types opt

[9] The present paper treats the quality index q as exogenous, with heterogeneous willingness-to-pay entering preferences additively via $\alpha_j q$. This differs from Veiga and Weyl (2016), who endogenise product design, showing that in selection markets insurers distort coverage characteristics away from the social optimum in response to heterogeneity in both risk and tastes, generating welfare losses beyond those identified in Rothschild-Stiglitz. In our framework competition operates through premiums for a given quality index rather than through joint price-quality design. Allowing quality to be endogenously chosen by insurers in a mixed public-private system, where the public offer itself pins down a quality benchmark, would be a natural extension: crowd-in at the intensive margin could then operate through product design adjustments as well as premium reductions, potentially amplifying the double-margin mechanism. For a related approach to Veiga and Weyl (2016), see Chade et al (2022).

[10] There can be four different types: Hh, Hl, Lh and Ll. We take two combinations of these types Lh and Hl and Ll and Hh as illustrative of the separating equilibrium under asymmetric information with averse selection when there is coexistence of public and private insurance.

for public coverage under adverse selection thus depends on the interaction of risk of illness type and willingness-to-pay, not on risk of illness alone.

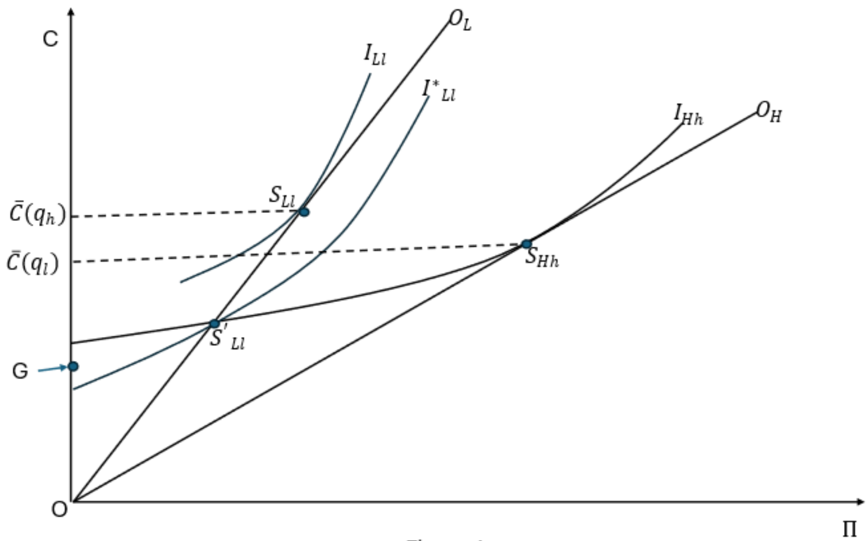

Figure 4

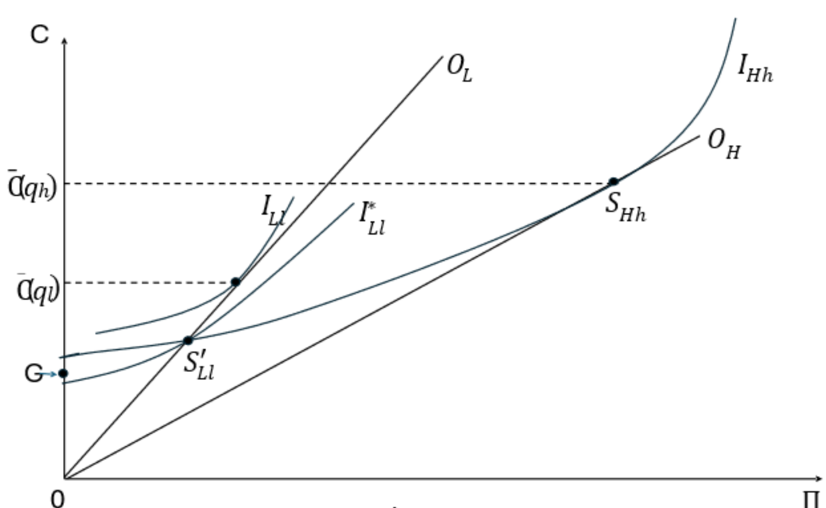

Figure 5

3.2 Advantageous Selection with coexisting public-private insurance with different preferences for quality

We now combine differences in risk aversion with quality heterogeneity, maintaining the negative correlation between risk aversion and illness probability, which motivates advantageous selection, and illustrate coverage movements by considering high risk of illness individuals with low willingness to pay for quality, Hl, types and low risk of illness individuals with high willingness to pay for quality, Lh, types. The equilibrium with public cover at G is partial pooling combined with this level of public coverage (P', G) rather than partial pooling combining with full private coverage at (P, S_{Hl}). Here some fraction ε' of Hl types choose public cover, and the remainder pool with Lh types at P', where

$$V_{Hl}(P') = V_{Hl}(G) \tag{10}$$

with the pooling premium satisfying

$$\Pi_{P'} \geq \frac{[\alpha\theta_L + (1-\alpha)(1-\varepsilon')\theta_H]}{\alpha + (1-\alpha)(1-\varepsilon')}\lambda_{P'}\bar{C} \tag{11}$$

as an equality and illustrated in Figure 6. The crowd-out of Hl types leaves a healthier private pool, improving contract terms at P' relative to P and generating crowd-in of marginal low-risk individuals, (Lh types), at the intensive margin who hold existing coverage through public insurance but crowd-in to the private market as they have a preference for quality.[11]

[11] We do not provide a formal proof in Appendix A of this intensive margin operation, but this result follows directly from the general proof provided for advantageous selection at Appendix A.2

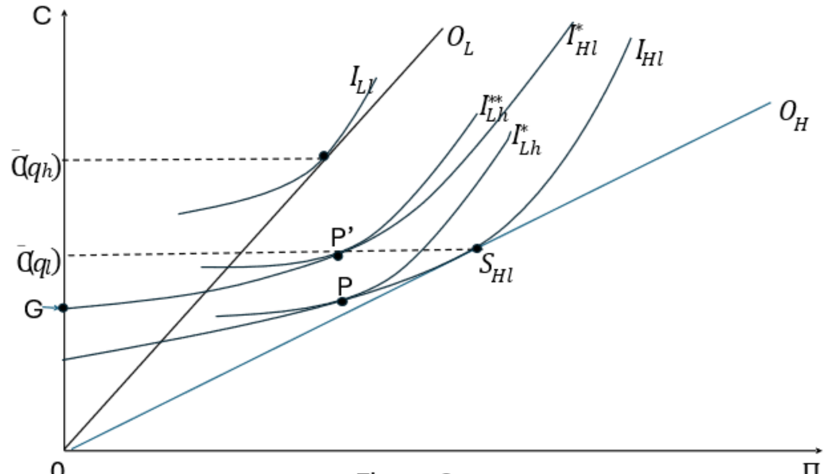

Figure 6

This partial pooling equilibrium, operating here through the intensive margin with quality heterogeneity, is the direct foundation for the double margin mechanism. As we now show, crowd-in can also operate at the extensive margin through the introduction of administrative costs and income heterogeneity.

4 Administrative Costs, Income Heterogeneity and the Double Margin Mechanism

The analysis so far has established the intensive margin effects of the double-margin mechanism where, under advantageous selection, crowd-out of high-risk types (Section 2.2) and of high-risk of illness individuals with low preference for quality (Section 3.2) improves the private contract terms at the partial pooling equilibrium, leading to intensive margin crowd-in. This section completes the double-margin mechanism by activating the extensive margin under a means-tested public insurance environment.

The key is actuarial unfairness induced by administrative costs. When private insurance is priced above its actuarial value, individuals who would have participated under fair pricing find insurance unattractive: their expected gain from coverage is positive but insufficient to cover the loading. We label these individuals M, a group of individuals who share the same illness probability and quality preference as high risk of illness individuals with low willingness to pay for quality (Hl), (such that $\theta_M = \theta_H$), but whose risk aversion is strictly lower. They are excluded from both markets, too rich for means-tested public insurance, too risk-tolerant to pay the loaded private premium. When crowd-out at the intensive margin lowers the pooling premium sufficiently, the M-type participation constraint is crossed and they crowd-in at the extensive margin. The three conditions, advantageous selection, administrative costs, and income-based eligibility, are each necessary for extensive margin crowd-in of the M types; together they are sufficient.[12]

Below we demonstrate the existence of the double-margin mechanism for a means-tested system analogous to Medicaid, then show how the institutional setting determines whether and how the mechanism operates under universal and mandatory coverage arrangements. These additional considerations of income heterogeneity and administrative costs create new margins of exclusion and deepen the selection dynamics already outlined.[13]

4.1 The Double-Margin Mechanism under Means-Tested Public Insurance

The double-margin mechanism operates through a single causal chain linking two otherwise independent margins. When public coverage improves, a larger fraction of high-risk of illness, low willingness to pay for quality (Hl) types exit the private pool, driving down the pooling premium through the zero-profit condition and eventually crossing the participation threshold of a group of individuals, labelled M (the same as Hl types but with lower risk aversion), who are excluded from both markets at the initial equilibrium: too wealthy for means-tested public coverage, too risk-tolerant to pay the loaded private premium. The key to this mechanism is actuarial unfairness induced by administrative costs ($k > 0$). When private insurance is priced above its actuarial value, the only route into coverage for M types is through the private market once crowd-out has reduced the premium sufficiently. We introduce administrative costs and an income distribution under which public coverage is means-tested, and show that these two additions, together with the advantageous selection established above, are jointly sufficient to generate the full double-margin mechanism.

To activate the extensive margin we introduce two further elements: administrative costs $k > 0$ that price private insurance above its actuarial value, and an income distribution

[12] The intensive margin crowd-in of Lh types established in Section 3.2 continues to operate simultaneously and its interaction with the M-type extensive margin crowd-in is discussed in Section 4.1.

[13] Throughout we use the following notation for pooling contracts. P denotes the pooling contract absent public insurance. P' denotes the pooling contract with public insurance and no administrative costs. P^* denotes the pooling contract with administrative costs prior to crowd-in of M-types. P^{**} denotes the pooling contract with administrative costs after crowd-in of M-types. Each successive contract reflects a further enrichment of the environment; the welfare ranking $P^{**} \succ P^* \succ P$ for L-types follows from the progressive improvement in the private risk pool.

$Y \sim [\underline{Y}, \bar{Y}]$ under which public coverage is restricted to individuals with income below threshold Y_b. Contract terms now satisfy

$$\Pi(q_j) \ \geq \ k + \theta_i \lambda_{i,j} \bar{C}(q_j), \qquad i = H, L;\ j = h, l \tag{12}$$

for group-specific contracts, and

$$\Pi_{P'} \ \geq \ k + \frac{\left[\alpha\theta_L + (1-\alpha)(1-\varepsilon')\theta_H\right]}{\alpha + (1-\alpha)(1-\varepsilon')}\, \lambda_P \bar{C} \tag{13}$$

for pooled contracts.[14]

4.1.1 The Double-Margin Mechanism Illustrated

Figure 7 presents the most general case: advantageous selection with risk heterogeneity, quality preferences, and administrative costs under means-tested public insurance. The figure introduces a group M of individuals who share the same illness probability and quality preference as Hl types ($\theta_M = \theta_H$), but whose risk aversion is strictly lower: their indifference curve I_M is less convex than I_{Hl} in premium–coverage space. Because their income exceeds the eligibility threshold ($Y_M > Y_b$), they cannot access public coverage at G regardless of its generosity. They are therefore excluded from both markets at the initial equilibrium: too risk-tolerant to pay the loaded private premium at P^*, and too wealthy to qualify for public coverage.

A chain of events now follows as the public offer G improves, a larger fraction ε' of Hl types exits the private pool at the intensive margin, improving private contract terms via the zero-profit condition. When the pooling premium falls far enough, the M-type participation constraint is crossed and these previously uninsured individuals enter the private market, the extensive margin. The two margins are therefore not independent: it is precisely the exit of Hl types that generates the fall in the premium that draws M types in.

Figure 7 systematically outlines this chain of events. Beginning at P^*, which represents the pooling contract prior to crowd-in of M types, the indifference curve I^*_M passes through P^* and intersects the horizontal axis to the left of k, confirming that at this premium M types prefer no insurance to private coverage. As the public offer G improves, ε' rises, Hl

[14] The comparative statics of (13) are straightforward. The administrative cost parameter k is exogenous; as it rises, premia rise and insurance becomes less attractive to all individuals. The proportion ε' is endogenous; as it rises, fewer Hl agents pool with L types and the pooling premium falls. Note that (13) describes the pooling contract P' prior to crowd-in of M types. Once M types enter the private market at P^{**}, the zero-profit condition takes the form of condition (b) in Definition 1, which keeps the Hl and M terms separate to reflect their distinct behavioural responses to G.

types exit, and the zero-profit condition drives the contract leftward and upward to P^{**}. At P^{**} the indifference curve I_M^{**} now passes above the origin. The improved contract terms have crossed the M-type participation threshold. The dashed curves I_{Lh}^* and I_{Lh}^{**} trace the corresponding improvement in contract terms for Lh types, illustrating the simultaneous intensive margin crowd-in (as noted in Section 4.1 below).

Note that because $\theta_M = \theta_H$, insurers cannot distinguish M types from Hl types on actuarial grounds.[15] No insurer can therefore profitably offer a contract designed to attract M types alone, which preserves the stability of the partial pooling equilibrium at P^{**} (established formally in Appendix A.3).

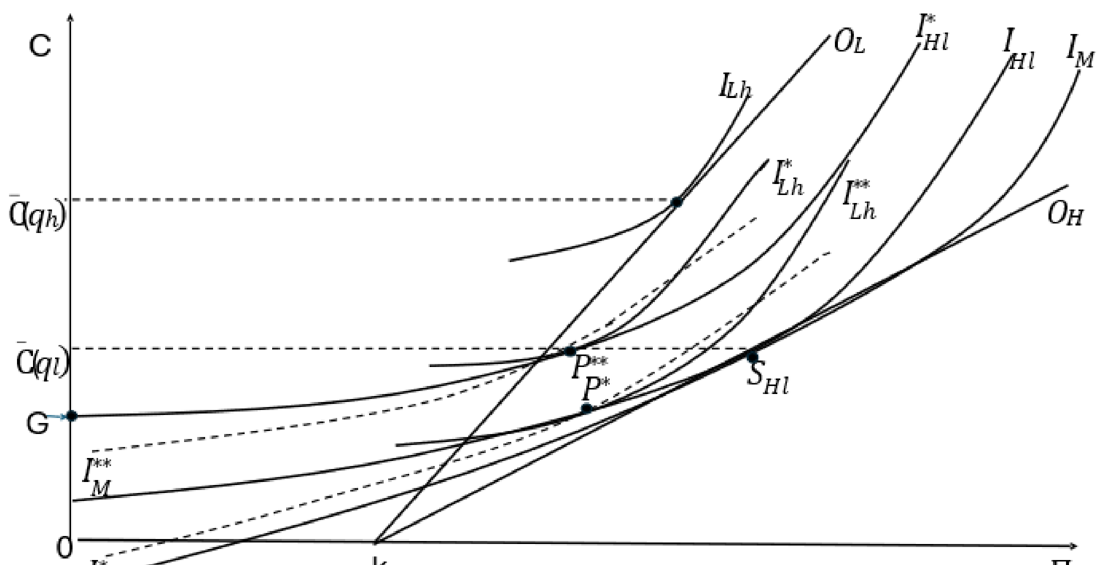

Figure 7

4.1.2 Formal Definition of the Double-Margin Mechanism

The three conditions identified above, (advantageous selection, administrative costs, and income-based eligibility), are necessary to evoke the external margin. Remove any one and a single margin operates. Moreover, crowd-in can only exist at the intensive margin if advantageous selection is present. Without advantageous selection the private premium rises

[15] See footnote 16 below

on crowd-out rather than falls; without administrative costs no individuals are priced out of the private market by premium loading; without income-based eligibility all risk types can access public coverage directly.

Definition 1 (Double-Margin Mechanism)**.** *In a mixed means-tested public–competitive private insurance system with administrative costs $k > 0$, income threshold Y_b, and advantageous selection, the double-margin mechanism operates through two causally linked margins.*

Margin 1 (Intensive margin crowd-out). A fraction ε' of Hl types with income $Y_{Hl} \leq Y_b$ switch from the private pooling contract P^{**} to public coverage G, where ε' is determined by the indifference condition

$$V_{Hl}(P^{**}) \;=\; V_{Hl}(G), \tag{a}$$

and P^{**} satisfies the zero-profit condition

$$\Pi_{P^{**}} \;=\; k + \frac{\alpha\theta_L + (1-\alpha)(1-\beta)(1-\varepsilon')\theta_H + (1-\alpha)\beta\theta_H}{\alpha + (1-\alpha)(1-\beta)(1-\varepsilon') + (1-\alpha)\beta}\,\lambda_{P^{**}}\bar{C}(q_h). \tag{b}$$

The three terms in the numerator and denominator correspond to the three groups in the private pool. Low-risk of illness types (L) (mass α) remain in the private pool throughout. High-risk of illness, low willingness to pay for quality (Hl) types (mass $(1-\alpha)(1-\beta)$) are present only in fraction $(1-\varepsilon')$, the remainder having taken public coverage at G; as G rises, ε' rises and their contribution to the pool shrinks. The M types, who are the same as Hl types but with lower risk aversion, (mass $(1-\alpha)\beta$) are locked into the private pool by the income eligibility condition $Y_M > Y_b$: they cannot access public coverage regardless of the generosity of G, so their mass enters condition (b) in full and is invariant to G. Since $\theta_M = \theta_H$, the Hl and M types share the same risk weight they may be combined as $(1-\alpha)\big[(1-\beta)(1-\varepsilon') + \beta\big]\theta_H$, but the two groups play distinct roles: only the Hl types responds to G through ε', while the M type do not.[16] The zero-profit condition (b) is therefore strictly decreasing in ε' alone, since $\theta_H > \theta_L$, and drives the pooling premium down as more Hl types exit.

[16] A number of supporting arguments are provided in Appendix A.3. The assumption $\theta_M = \theta_H$ is adopted for analytical tractability. The mechanism does not require this equality. For the proof of Lemma 2 in A.3 to hold, it is sufficient that M types have risk aversion low enough relative to their illness risk that the loaded private premium exceeds their reservation price, and income high enough to exclude them from means-tested public coverage. Both conditions are consistent with $\theta_M > \theta_H$. In that case the causal chain outlined below at (d) remains operative: the zero-profit condition (b) continues to be strictly decreasing in ε', and the M-type participation constraint in (c) is eventually crossed as $\Pi_{P^{**}}$ falls, though the threshold k^* is higher since a larger premium reduction is needed to make private coverage attractive to types whose actuarial risk exceeds that of Hl types. The equality $\theta_M = \theta_H$ is therefore a sufficient but not necessary condition for the mechanism; relaxing it raises the threshold but does not alter the qualitative results.

Margin 2 (Extensive margin crowd-in). Previously excluded M types, who are characterised as $\theta_M = \theta_H$, $Y_M > Y_b$, and risk aversion strictly below that of Hl types, satisfy $V_M(P^{**}) < V_M(0)$ at the pre-crowd-out premium. Once ε' rises sufficiently this reduces $\Pi_{P^{**}}$ below the M-type reservation premium,

$$V_M(P^{**}) \;\geq\; V_M(0). \tag{c}$$

The threshold k^* is the value of k at which condition (c) holds with equality at the equilibrium $\varepsilon'(G)$; crowd-in occurs if and only if $k > k^*$.

The causal chain linking the two margins is:

$$\partial G \uparrow \Rightarrow \; \partial\varepsilon' \uparrow \Rightarrow \; \partial\Pi_{P^{**}} \downarrow \Rightarrow \; \partial M_{\text{participation}} \uparrow . \tag{d}$$

Each link is strictly monotone, so the chain is robust to small perturbations. Crucially, the links are causally ordered rather than merely correlated, as each is triggered by the one before it, so that extensive margin crowd-in would not occur without the intensive margin crowd-out that initiates the chain.

Proposition 1 (Double-Margin Mechanism)**.** *In a mixed means-tested public–competitive private insurance system with administrative costs $k > 0$, income threshold Y_b, and advantageous selection, the double-margin mechanism of Definition 1 exists and constitutes a Nash equilibrium if and only if $k > k^*$, $\theta_M = \theta_H$ with strictly lower risk aversion for M than Hl types, and $Y_M > Y_b$. The three conditions are each necessary and jointly sufficient: without advantageous selection the private premium rises on crowd-out rather than falls; without administrative costs no individuals are priced out of the private market by premium loading; without income-based eligibility all risk types can access public coverage directly.*

Proof. The monotonicity of each link in chain (d) follows from Definition 1: $\partial\varepsilon'/\partial G > 0$ from condition (a); $\partial\Pi_{P^{**}}/\partial\varepsilon' < 0$ from condition (b) and $\theta_H > \theta_L$; and condition (c) is crossed if and only if $k > k^*$. Robustness to insurer deviation is established in Appendix A.3; the key step is Lemma 1, which shows that L types cross-subsidise Hl and M types in the pooling equilibrium, making profitable deviation impossible. □

Three further features of the mechanism deserve emphasis. First, the crowd-in at the extensive margin is not a coincidental by-product of crowd-out but is caused by crowd-out leading to an endogenous adjustment of private contract terms.[17] Note also that the

[17] This causal relationship operating across the public-private sectors distinguishes the double-margin mechanism from Encinosa (2003) who provides a taxonomy of crowd-out and crowd-in but no formal mechanism, and from Geruso et al. (2023) who identify intensive and extensive margin interactions in private markets with adverse selection but do not extend consideration to public-private insurance interactions.

mechanism is asymmetric with respect to the selection regime. Under adverse selection $\partial \Pi_{P^{**}} / \partial \varepsilon' > 0$, so crowd-out of low-risk types raises the private premium and chain (d) reverses, generating further exit rather than entry at the extensive margin. Advantageous selection is therefore not merely a technical assumption but the essential ingredient that makes crowd-in possible.

Second, crowd-out operates only through those high-risk of illness individuals with low willingness to pay for quality, (Hl types), who were previously buying private insurance and are now eligible for public coverage at G. High-risk individuals who could not afford private insurance initially exert no pressure on the private pool when public coverage expands, since they were not participating in the private market to begin with.

Third, while Proposition 1 focuses on the extensive margin, the intensive margin crowd-in of low-risk of illness, high willingness to pay (Lh) types established in Section 3.2 continues to operate simultaneously and plays a distinct role. As the pooling premium falls through chain (d), Lh types find private insurance increasingly attractive and their entry deepens the cross-subsidy of Lemma 1, allowing the zero-profit condition to sustain a lower premium and thereby reinforcing the stability of the equilibrium at P^{**}. Because this Lh type crowd-in is triggered by the same premium fall that draws M types in, the two crowd-in effects are mutually reinforcing. The Lh crowd-in is not, however, necessary to activate the extensive margin crowd-in of M types: Proposition 1 holds independently of any intensive margin crowd-in by Lh types, which helps stabilise the equilibrium once the mechanism is active rather than initiate it.

4.2 Extension to Universal Public Insurance Coverage Systems

The framework developed so far assumes individuals face a choice between no insurance, private insurance, and, if eligible, public insurance. While this accurately represents systems that exhibit means-tested public options, such as the U.S. Medicaid, many countries operate under different institutional arrangements that require reinterpretation of our model's baseline. In universal public insurance healthcare systems like the UK's NHS, the "no insurance" option at the origin (0,0) does not exist, as public coverage exists for all regardless of income level. We extend our framework to these settings by redefining the origin and contract space. In the UK NHS type system, all residents receive comprehensive healthcare funded through taxation, so there are no uninsured. We therefore re-label the y-axis as $\Delta\overline{C}(q_i)$ to indicate that this is supplemental private coverage, increasing in quality/coverage, over and above the public NHS coverage indicated as G at the origin and the origin is redefined as (0,G). The x-axis reflects the premium Π arising from supplemental private insurance cover, which incurs administration costs ($k > 0$). Under this reinterpretation, private insurance contracts represent either supplementary coverage (faster access, private rooms, choice of specialist) or substitutive coverage (opting out of NHS for specific treatments).

To draw out our institutional comparison we retain an M group type, as defined above, again assuming they have the same illness probability and preference for care as the Hl group. We again also assume that the M group have lower risk aversion than the Hl group, and consequently their indifference curve, I_M, is less convex than that of the Hl group. For example, income heterogeneity may define M-types who have the same probability of illness as high-risk types, $\theta_M = \theta_H$, but have differing risk aversion. Within the NHS setting, and given administration costs for private coverage, this translates into the whole of the M group being satisfied with the public sector NHS coverage. Supplementary private insurance coverage is unattractive to them; the marginal benefit of higher private coverage/quality is not justified given the private insurance premium for these less risk-averse types. As shown in Figure 8, they prefer public insurance at the origin, (0,G), to either of the contracts at P^{**} or S_{Hl}.

The double-margin mechanism outlined in Section 4.1 above requires an excluded group for the extensive margin to be activated. Under NHS-type universal coverage this group does not exist. As just stated, the M types are already publicly insured and have no reason to enter the supplementary private market in response to any premium reductions. The intensive margin crowd-out of a proportion of Hl types remains operative (defined as Margin 1 in Section 4.1). Now any improvement in the public offer draws a proportion (defined above as ε') of high risk of illness but low willingness to pay types, (Hl), out of the private supplementary private insurance market and consequently the private pool improves, meaning private insurers can lower premiums. Assuming advantageous selection remains present under NHS-type universal coverage, this crowd-out results in crowd-in of low risk of illness, high willingness to pay for quality types, (Lh), at the intensive margin. These are individuals who have NHS cover but have sufficiently strong quality preferences to find supplementary private insurance attractive once premiums fall. This intensive margin crowd-in mechanism is analogous in structure to that outlined at Section 3.2, although here private insurance is complementary to, rather than a substitute for, public coverage as reflected in the shift of the origin from $(0, 0)$ in Figure 6 to $(0, G)$ in Figure 8.[18]

[18]We note that Olivella and Vera-Hernández (2013) consider private insurance to only be a substitute in NHS-type institutional structures. This ignores the improvement in quality of coverage (such as avoiding waiting lists, dealing with senior clinicians only, private rooms, etc.) that private insurance coverage may offer over and above the public insurance offer in such institutional settings. In examining a model of competitive private insurance substituting for a universal public option, they derive testable predictions and find empirical evidence of adverse selection in the UK private health insurance market.

With non-excludable NHS-type public coverage what is absent is the extensive margin crowd-in of M types. Since all individuals including M types hold public coverage at the origin $(0, G)$, there is no excluded group whose participation constraint can be crossed by a premium reduction. Therefore, under an NHS arrangement, the mechanism reduces to an intensive margin, with private supplementary insurance serving predominantly low-risk, high quality-preference types (Lh) together with the residual fraction $(1 - \varepsilon')$ of Hl types who do not exit to public coverage.

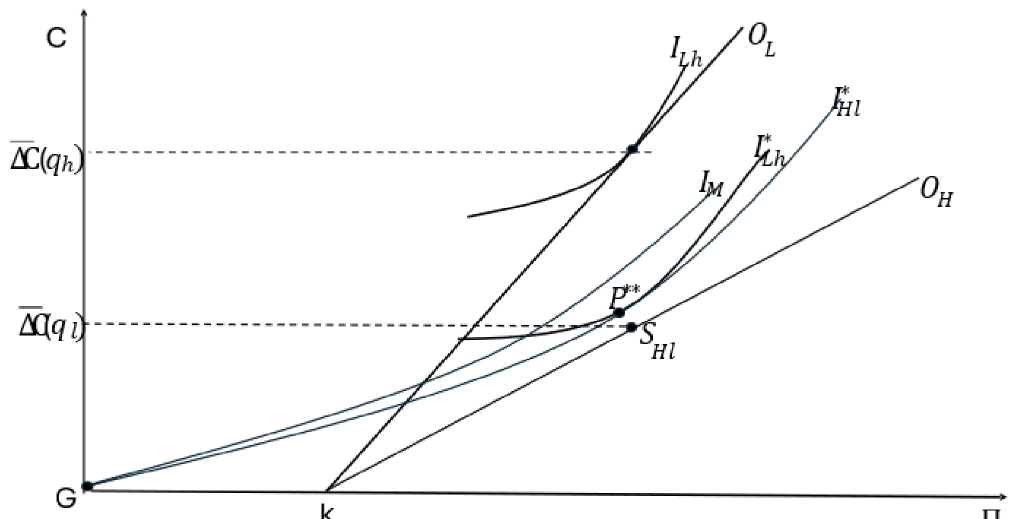

Figure 8

4.3 Extension to Mandatory Public Insurance Coverage Systems

We now extend the analysis to a system of mandatory basic insurance with voluntary supplementary coverage, analogous to the Dutch system in which a regulated basic package is mandated by law and insurers compete on price and service quality and through supplementary coverage offered on a voluntary basis.[19] The double-margin logic carries over to this setting, but the analysis is more complex. A deductible now operates across multiple intensive and extensive margins simultaneously, and the interaction among them cannot be defined analytically without empirical input on the joint distribution of illness risk, income, and risk aversion. To consider the structure of the problem we can redefine the contract

[19]For details of the Dutch health insurance system see van Kleef et al. (2020). One might also characterise the US Medicare and Medigap system in these terms.

space (not drawn) as follows. The Point $(0, G)$ now represents mandatory basic insurance at community-rated premium. The origin is thus not zero coverage but the mandated base coverage. The individual now faces two premium obligations simultaneously: a mandatory basic premium Π_{basic} that cannot be avoided, and a voluntary supplementary premium Π_{supp} chosen optimally. The contract space is defined over $(\Delta\overline{C}(q_i), \Pi)$, representing the range of supplementary coverage choices available above the mandated base. The deductible D creates a non-linearity in out-of-pocket costs: the individual bears the full cost of illness up to D, after which insurance covers its contracted share of the remainder. Incorporating these features, preferences take the form:

$$\begin{aligned} V_i = {} & \theta_i U_i \left(Y_i - \min(C, D) - \max(0, C - D)\left(1 - \rho_{\text{basic}} - \rho_{\text{supp}}\right) - \Pi_{\text{basic}} - \Pi_{\text{supp}}\right) \\ & + (1 - \theta_i)\, U_i \left(Y_i - \Pi_{\text{basic}} - \Pi_{\text{supp}}\right) \end{aligned} \tag{14}$$

where D represents deductible payments and ρ_{basic} , ρ_{supp} represent basic and supplementary coverage levels respectively. The individual's payment structure becomes complicated as: $\min(C, D)$ represents out-of-pocket deductible payment (capped at total cost C if $C < D$); $\max(0, C - D)$ are costs above the deductible; $(1 - \rho_{\text{basic}} - \rho_{\text{supp}})$ is the portion not covered by insurance after the deductible is met. The individual bears the full cost up to D, then insurance covers its portion. [20]

In the Netherlands the mandatory basic tier operates under a risk equalisation scheme administered by the National Health Authority (NZa), under which insurers receive risk-adjusted transfers $\tau(\theta_i)$ that modify the zero-profit condition for basic insurance to

$$\Pi_{basic} \geq k + \theta_i \lambda_{basic} \bar{C} - \tau(\theta_i) \tag{15}$$

For completeness we retain the M types, who have the characteristics of low risk aversion relative to their risk of illness. However, a sufficiently generous equalisation scheme compresses the actuarial loading and shrinks the excluded M group at the basic tier. If equalisation is actuarially complete the mechanism that generates M types is eliminated entirely. Critically, equalisation applies only to the basic tier. The voluntary, supplementary market retains the full loading structure of equations (13) and (14), and it is at this tier that the double margin mechanism, if operative, will be located in the Dutch setting. This asymmetry between tiers is the primary reason why the precise equilibrium cannot be resolved theoretically without empirical input on equalisation generosity and the distribution

[20] One natural extension would allow individuals to choose their deductible level within the basic tier, generating further heterogeneity in effective coverage; we leave this for future work.

of illness risk, income, and risk aversion.

This creates a complex selection dynamic. Low-risk individuals may choose high deductibles and skip any supplementary coverage. High-risk individuals face a trade-off: supplementary coverage helps with costs above the deductible but doesn't reduce the deductible burden. The deductible creates a 'first-cost' exposure that affects both basic and voluntary supplementary coverage value with what we now refer to as the extensive margin in this setting applying to the purchase of voluntary supplementary coverage levels beyond the mandated cover. We refer to the intensive margin in this setting as operating at both the mandated and supplemental coverage as a deductible is charged on both. The application of a deductible on both the mandated and supplementary level will mean that single-crossing (with adverse selection) and double-crossing (with advantageous) properties could apply to both mandated and voluntary supplementary insurance coverage. In other words, there could be differential responses by low-risk and high-risk individuals at different income levels. Sorting of the different types into the different levels of supplementary coverage therefore becomes complex and conditioned on risk-type, levels of risk aversion (and income levels), and preferences for coverage and treatment quality.

The double-margin mechanism can exist again, but requires further modification as complexities are introduced due to the impact of deductibles across the range of contracts, creating multiple margins. While our model cannot, without further amendment informed by empirical observation, give a full characterisation of the potential outcomes in this institutional setting, we give the following example. A modified crowd-out will exist where improvements in public sector coverage leads to less purchasing of supplemental insurance coverage by high-risk individuals with low willingness to pay for quality (Hl types). This in turn leads to modified crowding-in where, depending on the relative size of the deductibles and relative coverage levels of the basic mandated insurance compared to the private supplemental insurance, various types will find the adjusted private insurance contracts, amended as a result of the crowding-out, attractive. The crowd-in effect at the extensive margin in this institutional setting, if operative, runs through M types entering the voluntary supplementary market once crowd-out of Hl types has reduced the supplementary premium sufficiently and not through low risk of illness with high willingness to pay for quality types, (Lh), who are already participating at the basic tier and are not marginal. Whether this M-type crowd-in occurs depends on the residual loading at the supplementary tier after risk equalisation has set the basic tier and on whether a population of individuals of M types, (with the characteristics of low risk aversion relative to their illness risk), who are ineligible for additional public subsidy exists at the supplementary margin. This is an empirical question that the theoretical framework alone cannot resolve. However, this example broadly aligns with empirical observation from the Netherlands, where supplementary insurance uptake correlates with income and education suggesting advantageous selection in the supplementary market and, although crowd-in is never referred to in the studies, is indicative of crowd-in in this institutional setting (Alessie et al., 2020; Handel et al., 2024; Remmerswaal et al., 2023).[21]

4.4 Institutional Dependence of the Double Margin Mechanism

Proposition 1 established the double-margin mechanism for the means-tested case. We now state a second main result, Proposition 2, which concerns how the institutional setting determines whether and how the mechanism operates.

Proposition 2 (Institutional Dependence)**.** *The double-margin mechanism is fully active only under means-tested public insurance. Under NHS-type universal coverage the extensive margin does not exist and only intensive margin crowd-out and crowd-in occurs. Under mandatory basic insurance with supplementary private coverage the mechanism operates in modified form across multiple margins, and the precise equilibrium is an empirical matter.*

Proof. As established in Proposition 1, the double-margin mechanism operates fully under a means-tested public insurance system. Under universal coverage, M types are already publicly insured, so the condition $Y_M > \bar{Y}_b$ required for M-type ineligibility in Definition 1 cannot hold; the extensive margin is therefore absent and the feedback chain operates through intensive margin crowd-in. Under mandatory basic insurance, two features modify the mechanism. First, risk equalisation transfers at the basic tier, as in equation (18), compresses the effective loading and reduces the size of any excluded M group there, potentially to zero if equalisation is near-complete. Second, since equalisation does not extend to the voluntary supplementary tier, the loading structure that generates M types is preserved in the supplementary market. The double-margin logic therefore applies at that tier, with the mandatory basic contract replacing the no-insurance origin. Together these features imply that the mechanism remains operative but spans multiple extensive and intensive margins simulataneously. The precise equilibrium attained depends on equalisation generosity, deductible levels, and the distribution of risk aversion and income, and is therefore an empirical matter[22] □

[21] Two empirical questions follow directly. First, how complete is NZa equalisation at the basic tier? If near-complete, the double margin mechanism is confined to the supplementary market and identification of effects should focus there. Second, what is the size and composition of the residual non-participating (in the supplemental market) population? If this group shares the characteristics of M types in the means-tested framework, a modified extensive margin exists even at the basic tier. Resolving these questions using Dutch administrative data, which record equalisation transfers directly, is the natural next step.

[22] In practice, mandatory participation is imperfectly enforced in some jurisdictions, leaving a residual uninsured population whose outside option differs from the mandated contract and who may constitute a

The institutional comparison reveals that the same underlying selection dynamics built around advantageous selection, administrative loading and income heterogeneity, generate qualitatively distinct equilibrium outcomes depending solely on whether and how public eligibility is restricted. It follows that applied analysis cannot treat institutional arrangements as interchangeable. Standard analysis, which captures only the intensive margin crowd-out, will systematically misstate the welfare consequences of public insurance expansions whenever intensive margin crowd-in and/or extensive margin crowd-in activating a double-margin is present. When the double-margin is present welfare effects will be monotonic (as we establish formally in Appendix B).

5 Conclusion

This paper has developed a new theoretical account of how public and private health insurance interact, centred on a double-margin mechanism, operating on both the intensive and extensive margins of insurance consumption. The dominant view in the literature treats public insurance expansion as a substitute for private coverage, generating crowd-out at the intensive margin as individuals shift from private to public provision. We have shown this account to be systematically incomplete. Under advantageous selection, crowd-out of high-risk of illness, low quality-preference individuals (Hl types) from the private insurance pool endogenously improves private contract terms, generating crowd-in at the intensive margin as low-risk, high quality-preference individuals (Lh types) find private insurance increasingly attractive. When public insurance is means-tested and private insurance is priced above its actuarial value due to administrative costs, a further and distinct crowd-in mechanism is activated at the extensive margin: previously excluded individuals who have the same characteristics as Hl types, but additionally exhibit moderate risk aversion and have sufficient income (which we have defined as M types), who were previously too wealthy for public coverage and too risk-tolerant to pay the initial loaded private premium, enter the private market once the private pooling premium falls sufficiently. We show that these two margins are not independent effects that happen to coincide, but are causally linked through a single chain of events. An improvement in public generosity raises the fraction of Hl types exiting the private pool, which lowers the pooling premium through the zero-profit condition, which eventually crosses the participation threshold of M types. We establish (Proposition 1) that three conditions are necessary and jointly sufficient for this mechanism: advantageous selection, administrative loading, and income-based eligibility for public coverage. Remove any

modified M group at the basic extensive margin. A full characterisation of this case requires empirical input on enforcement rates and is left for future work.

one and the mechanism collapses to a single intensive margin.

We also establish (Proposition 2) that the institutional setting matters in the operation of the mechanism. Under means-tested public insurance, analogous to Medicaid in the United States, the double-margin operates in full. Under NHS-type universal coverage the extensive margin is absent by construction, as M types are already publicly insured and no excluded population exists whose participation constraint can be crossed by a premium reduction. Private supplementary insurance ultimately serves predominantly low-risk, high quality-preference individuals alongside the residual fraction of Hl types remaining in the private pool. Under mandatory basic insurance with voluntary supplementary coverage, as in the Netherlands, deductibles operate across multiple tiers simultaneously and generate a more complex selection dynamic. The double-margin logic carries over in modified form, but the precise equilibrium is not fully characterised without empirical input on the generosity of risk equalisation, deductible levels, and the joint distribution of income, illness risk, and risk aversion.

The welfare implications follow directly, which we detail in Appendix B. Under an expansion of public generosity with eligibility held fixed, the feedback chain we outline (in Definition 1) drives the pooling premium down, improving contract terms for all groups remaining in the private market. Abstracting from the fiscal cost of the expansion, the welfare gain is unambiguously Pareto improving (formally established in Appendix B.2.1). Under a reduction in the eligibility threshold to admit M types to public coverage, the welfare effects are more nuanced. The fiscal cost of extending public eligibility must be weighed against the combined welfare gains to all groups from the improved private risk pool and from M types accessing coverage. We establish (in Appendix B.2.2, eq. B.10), that this condition is more likely to hold when the M group is small relative to the high-risk population.

These results reframe the policy question. Standard crowd-out analysis, even when extended to capture intensive margin crowd-in, will systematically misstate the welfare consequences of public insurance expansions whenever the double-margin is active, as it omits the extensive margin crowd-in of previously uninsured individuals and the endogenous improvement in private insurance contract terms that drives it. The appropriate empirical strategy requires data that jointly record individual insurance status, income relative to the eligibility threshold, and private premium movements over time. Such data exist in both the US Medicaid context and in the Dutch system. Constructing credible estimates of the double-margin mechanism, its size, its distributional incidence, and whether the extensive margin crowd-in is sufficient to offset the fiscal cost of expanded eligibility, is the natural next step and is left for future work.

A Equilibrium Characterisation under Adverse and Advantageous Selection

A.1 The Separating Equilibrium under Adverse Selection with Public Option (G)

With the introduction of a public insurance option at zero premium providing coverage G, the agent's choice set expands from $\{S_H, S'_L\}$ to $\{G, S_H, S'_L\}$. The public option yields expected utility

$$V_i(G) = \theta_i U_i\big(Y - \bar{C} + G\big) + (1-\theta_i)U_i(Y), \quad i = H, L, \tag{A.1}$$

and is available to all individuals regardless of type. We retain the single-crossing property throughout, so that H-type indifference curves are flatter than L-type indifference curves in premium–coverage space.

Equilibrium conditional on G. Three cases arise depending on the generosity of G.

Case 1 (G insufficiently attractive). If $V_i(G) < V_i(S_i)$ for both $i = H, L$, neither type finds public coverage attractive. The standard Rothschild–Stiglitz separating equilibrium (S_H, S'_L) obtains unchanged.

Case 2 (G moderately generous). When $V_H(G) \geq V_H(S_H)$ but $V_L(G) < V_L(S'_L)$, high-risk types prefer public coverage and exit the private market. The private pool contracts to L types only, who obtain full coverage at the actuarially fair premium $\Pi_L = \theta_L \bar{C}$.

Case 3 (G sufficiently generous to attract L types — Figure 2). At level $G_2 > G_1$, equation (5) of the main text gives

$$V_L(G_2) \geq V_L(S'_L)\,, \qquad V_H(G_2) < V_H(S_H). \tag{A.2}$$

Under asymmetric information and single crossing, L types are attracted to public coverage before H types. The deductible $(1-\lambda_L)\bar{C}$ embedded in S'_L—required to satisfy the H-type incentive-compatibility constraint—makes private coverage sufficiently unattractive that L types prefer the public bundle at $(0, G_2)$. Case 3 is the equilibrium illustrated in Figure 2; we verify it formally below.

Verification of Case 3. The proposed equilibrium has L types choosing G and H types choosing $S_H = (\Pi_H, \bar{C})$ with $\Pi_H = \theta_H \bar{C}$. Four conditions must hold.

1. *H-type participation.* H types prefer S_H to public coverage: $V_H(S_H) > V_H(G_2)$, which holds by the definition of Case 3.

2. *L-type participation.* L types prefer G to any private contract. Since S_L' was the most attractive private contract for L types absent G, it suffices that

$$V_L(G_2) \geq V_L(S_L') > V_L(S_H).$$

   The first inequality holds by the Case 3 condition; the second follows from the L-type participation constraint in the original Rothschild–Stiglitz equilibrium.

3. *H-type incentive compatibility.* With L types absent from the private market, H types have no incentive to mimic L. The only relevant deviation is to switch from S_H to G, which is precluded by condition (i).

4. *Zero-profit condition.* With only H types in the private pool, $\Pi_H = \theta_H \bar{C}$ holds as an equality. No cross-subsidy arises and the insurer breaks even.

All four conditions are satisfied, confirming that Case 3 constitutes a Nash equilibrium. The key insight is that the introduction of G does not alter the single-crossing logic that prevents pooling; it alters only which type exits the private market first, and this depends entirely on the generosity of G relative to the deductible embedded in the separating contract S_L'.

Non-existence of pooling equilibria. We show that no pooling contract can be sustained as a Nash equilibrium, whether or not a public option G is present. Suppose, contrary to this claim, that a pooling contract $P = (\Pi_P, \lambda_P \bar{C})$ exists in which both H and L types participate. The zero-profit condition requires

$$\Pi_P = [\alpha\theta_L + (1-\alpha)\theta_H]\, \lambda_P \bar{C}, \tag{A.3}$$

where α is the proportion of L types. Since $\theta_H > \theta_L$, the pooling premium exceeds the actuarially fair premium for L types: $\Pi_P > \theta_L \lambda_P \bar{C}$.

By the single-crossing property, the L-type indifference curve through P is steeper than the H-type indifference curve through P in premium–coverage space. A deviating insurer can therefore offer a contract $P^* = (\Pi^*, \lambda^* \bar{C})$ with $\lambda^* < \lambda_P$ and $\Pi^* = \theta_L \lambda^* \bar{C}$ that lies above the L-type indifference curve through P but below the H-type indifference curve through P. Such a contract attracts L types but not H types, satisfies the zero-profit condition for a pool of L types alone, and is therefore strictly profitable. The pooling contract P is thus vulnerable to profitable deviation and cannot constitute a Nash equilibrium.

This argument is unaffected by the presence of the public option G. In Cases 2 and 3 above, H types or L types have already exited to public coverage, so the relevant private pool is smaller but homogeneous; a fortiori no pooling deviation is profitable. In Case 1 the public option is not binding and the standard Rothschild–Stiglitz non-existence result applies directly. The separating equilibrium is therefore the unique Nash equilibrium across all three cases.

A.2 The Partial Pooling Equilibrium under Advantageous Selection with Public Option G

We now assume asymmetric information under advantageous selection: L types are more risk-averse than H types, so the negative correlation between risk aversion and illness probability causes indifference curves to cross twice in premium–coverage space (the double-crossing property). The introduction of public option G at zero premium augments the agent's choice set to $\{G, P\}$, where $P = (\Pi_P, \lambda_P \bar{C})$ is the private pooling contract. We show that a partial pooling equilibrium with public option G constitutes a Nash equilibrium, and that the double-crossing property is what prevents Rothschild–Stiglitz unravelling.

The double-crossing property and stability. Under advantageous selection, L-type indifference curves are strictly more convex than those of H types and cross the H-type indifference curve exactly twice. At any pooling contract P in the relevant region, the L-type indifference curve lies above the H-type indifference curve both to the left and to the right of P, with tangency at P itself. This double-crossing property has two consequences that are critical to the proof below.

First, it rules out the Rothschild–Stiglitz unravelling that afflicts the adverse selection case. Under single crossing any pooling contract can be profitably undercut by a contract tailored to low-risk types; under double crossing, the L-type indifference curve lying above the H-type curve on both sides of P means that any contract more attractive to L types than P is also more attractive to H types, so no insurer can profitably separate the pool.

Second, there is a unique pooling contract consistent with zero profit for any given ε, the fraction of H types choosing public coverage. Each value of ε maps to a unique pooling premium through the zero-profit condition, and by the strict convexity of indifference curves there is a unique ε^* satisfying the H-type indifference condition $V_H(P') = V_H(G)$. This pins down the equilibrium.

Proposed equilibrium with public option G. A fraction ε' of H types choose public coverage G; the remaining fraction $(1-\varepsilon')$ pool with all L types at the private contract $P' = (\Pi_{P'}, \lambda_{P'} \bar{C})$.

The equilibrium objects satisfy:

$$V_H(P') = V_H(G), \tag{A.4}$$

$$\Pi_{P'} = \frac{\alpha\theta_L + (1-\alpha)(1-\varepsilon')\theta_H}{\alpha + (1-\alpha)(1-\varepsilon')}\lambda_{P'}\bar{C}, \tag{A.5}$$

$$V_L(P') > V_L(G). \tag{A.6}$$

Condition (A.4) is the indifference condition for the marginal H type; it determines ε' jointly with the zero-profit condition (A.5). Condition (A.6) states that all L types strictly prefer the private pooling contract to public coverage, so no L types exit to G.

Verification. Since $\partial V_H(G)/\partial G > 0$ and $\partial V_H(P')/\partial\varepsilon' > 0$ (a larger ε' improves the private pool and raises $V_H(P')$ via the zero-profit condition), condition (A.4) implicitly defines $\varepsilon' = \varepsilon'(G)$ with $\partial\varepsilon'/\partial G > 0$: a more generous public offer draws more H types out of the private pool. The zero-profit condition (A.5) is strictly decreasing in ε' since $\theta_H > \theta_L$, so higher ε' lowers the pooling premium and improves contract terms at P'. By the double-crossing property, L types' indifference curves are strictly above H types' curves at P'; since H types are indifferent between P' and G, L types strictly prefer P' to G, confirming (A.6).

Robustness to insurer deviations. We verify that no insurer can profitably deviate from this equilibrium. Let $S^* = (\Pi^*, \lambda^*\bar{C})$ be any alternative contract.

Deviation targeting L types only. For S^* to attract L types it must satisfy $V_L(S^*) > V_L(P')$. By the double-crossing property, any contract lying above the L-type indifference curve through P' also lies above the H-type indifference curve through P'. Since H types at P' satisfy $V_H(P') = V_H(G)$, any contract preferred by L types to P' also satisfies $V_H(S^*) > V_H(G)$, attracting marginal H types as well. Including H types raises the zero-profit premium above $\Pi_{P'}$, making the deviation unprofitable.

Deviation targeting H types only. Any contract designed to attract only the $(1-\varepsilon')$ fraction of H types pooling at P' must charge at least the actuarially fair premium $\theta_H\lambda^*\bar{C}$. Since $\theta_H > \theta_L$, this exceeds the pooling premium $\Pi_{P'}$ for any given coverage level (by the cross-subsidy embedded in the pooling contract). H types at P' are already indifferent between P' and G; any contract requiring a strictly higher premium for the same coverage satisfies $V_H(S^*) < V_H(P') = V_H(G)$, so H types prefer G to the deviation. No profitable deviation targeting H types alone exists.

Deviation offering a new pooling contract. Consider an alternative pooling contract S^{**} covering both types. For S^{**} to attract L types it must offer better terms than P'. Either improvement raises the zero-profit premium for the deviating pool above $\Pi_{P'}$ (since

H types, who are more costly, are also attracted by double-crossing), violating zero-profit. At coverage $\lambda^{**} = \lambda_{P'}$, any premium $\Pi^{**} < \Pi_{P'}$ attracts both types but earns negative expected profit since the pooled expected loss exceeds Π^{**}. No profitable pooling deviation exists.

All three deviation cases fail, confirming that the partial pooling equilibrium (P', G) constitutes a Nash equilibrium. This equilibrium is illustrated in Figure 3. The endogenous improvement in private contract terms as ε' rises with G constitutes the intensive margin of the double-margin mechanism, completed formally in Appendix A.3.

A.3 Proofs of Proposition 1 and Supporting Lemmas

The proof establishes that the double-margin equilibrium defined in Definition 1 constitutes a Nash equilibrium. The argument proceeds in three steps. Lemma 1 establishes a key property of the pooling premium, that low-risk types cross-subsidise high-risk types, which is used in the deviation analysis. Lemma 2 then shows that no insurer can profitably attract only Hl and M types away from the equilibrium, and Lemma 3 shows that deviations targeting the L group are equally unprofitable. Together these lemmas confirm that the equilibrium is robust to all possible deviations, completing the proof of Proposition 1.

Model Environment. We consider the means-tested market institution with three types of agents:

- Hl types (mass $(1-\alpha)(1-\beta)$): high probability of illness, low willingness to pay for quality, highly risk-averse.
- L types (mass α): low probability of illness, high willingness to pay for quality, moderately risk-averse.
- M types (mass $(1-\alpha)\beta$): same illness probability as Hl types ($\theta_M = \theta_H$), same quality preference, but strictly lower risk aversion, and ineligible for public insurance since $Y_M > \bar{Y}_b$.

All population shares sum to one: $\alpha + (1-\alpha)(1-\beta) + (1-\alpha)\beta = 1$. Agents choose between public insurance (zero premium, baseline coverage G) and private insurance (premium Π, coverage $\lambda_p \bar{C}$, subject to administrative cost k). The timing is: (1) public insurance terms are exogenously set; (2) private insurers offer contracts subject to zero-profit and administrative cost k; (3) agents choose between public and private insurance; (4) risk is realised and payoffs determined.

Equilibrium Definition. An equilibrium consists of a private contract $S = (\Pi_P, \lambda_P \bar{C})$ satisfying the zero-profit condition

$$\Pi_P = \frac{\alpha\theta_L + (1-\alpha)(1-\beta)(1-\varepsilon')\theta_H + (1-\alpha)\beta\theta_H}{\alpha + (1-\alpha)(1-\beta)(1-\varepsilon') + (1-\alpha)\beta} \lambda_p \bar{C} + k, \tag{A.7}$$

with the indifference condition for the marginal Hl type

$$V_{Hl}(\Pi_P, \lambda_p \bar{C}) = V_{Hl}(G), \tag{A.8}$$

and the initial partition of agents satisfying

$$V_{Hl}(\text{private}) = V_{Hl}(\text{public}), \tag{A.9}$$

$$V_L(\text{private}) > V_L(\text{public}), \tag{A.10}$$

$$V_M(\text{private}) < V_M(\text{no insurance}). \tag{A.11}$$

The three terms in the numerator of (A.7) correspond to the three groups in the private pool. L types (mass α) remain throughout. Hl types contribute only their remaining fraction $(1-\varepsilon')$, as a fraction ε' have taken public coverage. M types (mass $(1-\alpha)\beta$) are locked into the private pool by the income eligibility condition and enter in full. Since $\theta_M = \theta_H$, the Hl and M terms share the same risk weight; however only the Hl term responds to G through ε', while the M term is invariant to G.

As public insurance improves, ε' rises, condition (A.7) implies a lower premium, and condition (A.11) eventually reverses so that M types enter the private market. This is the double-margin mechanism. The equilibrium system determining ε' and M-type entry simultaneously is:

1. $V_{Hl}(\Pi_P, \lambda_p \bar{C}) = V_{Hl}(G)$ [indifference condition]
2. $\Pi_P = \mathbb{E}[\theta_i \mid i \in \text{private pool}]\lambda_P \bar{C} + k$ [zero profit]
3. $V_M(\Pi_P, \lambda_p \bar{C}) \geq V_M(0)$ [participation of M types]
4. Expected risk $= f(\varepsilon', \text{entry of } M \text{ types})$ [risk pool composition]

Proof of Robustness to Deviation. We now verify that no insurer can profitably deviate from this equilibrium. The proof uses three lemmas.

Lemma 1. *In the pooling equilibrium, L types cross-subsidise Hl and M types.*

Proof. From the zero-profit condition (A.7),

$$\Pi_P = \frac{\alpha\theta_L + (1-\alpha)(1-\beta)(1-\varepsilon')\theta_H + (1-\alpha)\beta\theta_H}{\alpha + (1-\alpha)(1-\beta)(1-\varepsilon') + (1-\alpha)\beta}\,\lambda_p \bar{C} + k. \tag{A.12}$$

Since $\alpha > 0$ and $0 < \theta_L < \theta_H = \theta_M$, the weighted average risk in (A.12) lies strictly below θ_H, so the pooling premium is below the actuarially fair premium for Hl and M types. L types therefore cross-subsidise both Hl and M types in the pool. Note that although Hl and M types carry identical actuarial risk θ_H, they differ in risk aversion and eligibility: Hl types may exit to public coverage as G rises, reducing their mass in the pool, while M types cannot and remain at full mass $(1-\alpha)\beta$ throughout. □

This cross-subsidy is what makes the pooling contract attractive to Hl types and, once the premium falls sufficiently, to M types. It also underlies the deviation analysis: any insurer attempting to attract only Hl and M types must charge at least the actuarially fair premium for those types, which by Lemma 1 exceeds the pooling premium and therefore makes the deviation unattractive to Hl types who are already indifferent at the pooling contract. Why the same deviation is unattractive to M types, whose lower risk aversion generates a different participation constraint, is established in Lemma 2 below.

Lemma 2. *No insurer can profitably offer a contract $S^* = (\Pi^*, \lambda^* \bar{C})$ that attracts only Hl and M types away from the equilibrium.*

Proof. For S^* to attract Hl and M types it must satisfy their participation constraints:

$$V_{Hl}(S^*) \geq V_{Hl}(G), \tag{A7}$$

$$V_M(S^*) \geq V_M(0). \tag{A8}$$

For the deviation to be profitable it must satisfy the zero-profit constraint, noting $\theta_{Hl} = \theta_M = \theta_H$:

$$\Pi^* = \theta_H \lambda^* \bar{C} + k. \tag{A9}$$

Case 1: $\lambda^* \leq \lambda_P$. From the equilibrium conditions, $V_{Hl}(\text{equilibrium}) = V_{Hl}(G)$ and $V_M(\text{equilibrium}) = V_M(0)$. Since M types are less risk-averse than Hl types, for any contract with $\lambda^* \leq \lambda_P$: if $V_{Hl}(S^*) \geq V_{Hl}(G)$, then $V_M(S^*) < V_M(\text{equilibrium}) = V_M(0)$, violating (A8).

Case 2: $\lambda^* > \lambda_P$. The zero-profit constraint (A9) requires $\Pi^* \geq \theta_H \lambda^* \bar{C} + k$. By Lemma 1 the equilibrium premium satisfies the pooling condition, so $\Pi^* > \Pi_P$ for identical coverage levels. Since Hl types are already indifferent between the equilibrium contract and public coverage G, any contract requiring a higher premium cannot satisfy (A7).

In both cases no deviation can simultaneously satisfy (A7) and (A8) while remaining profitable. The key tension is that Hl and M types carry identical actuarial risk θ_H but differ in risk aversion: contracts attractive to the more risk-averse Hl types are unattractive to M types, and vice versa, while both require the same actuarially fair premium. □

Lemma 3. *Deviations targeting the L group, or any combination including the L group, are unprofitable.*

Proof. Any deviation targeting only L types would also attract Hl types by the double-crossing property established in Appendix A.2, since a contract more attractive to L types than the equilibrium pooling contract lies inside the region where both types' indifference curves overlap. Including Hl types violates the zero-profit condition. Deviations targeting L and M types would similarly attract Hl types since $\theta_{Hl} = \theta_M$ means insurers cannot distinguish between them on risk grounds. Deviations targeting all types are either equivalent to the equilibrium offer or unprofitable by the zero-profit condition. □

Lemmas 1–3 jointly establish that no insurer can profitably deviate from the double-margin equilibrium. The equilibrium is therefore a Nash equilibrium, completing the proof of Proposition 1. □

B Welfare Analysis

This appendix examines the welfare implications of the equilibria characterised in the main text, complementing the existence and institutional dependence results of Propositions 1 and 2. The central question is whether expanding public insurance can be welfare improving, and under what conditions. Section B.1 analyses the adverse selection case of Section 2.1, identifying under-insurance of L types as the source of deadweight loss in the separating equilibrium, and showing that mandated minimum coverage can generate a Pareto improvement when the proportion of low-risk types exceeds a critical threshold α^* at which the efficiency gains from reducing under-insurance are sufficient to finance the cross-subsidy required to compensate H types. Section B.2 analyses welfare under the double-margin equilibrium of Proposition 1, distinguishing two policy experiments with qualitatively different properties. The first is an expansion of public coverage G with eligibility held fixed: because the feedback chain of Definition 1 drives the pooling premium down, all groups benefit and the welfare gain is unambiguously Pareto improving, abstracting from the fiscal cost of G. The second is a reduction in the eligibility threshold to admit M types to public coverage: this removes M types from the private pool, raises the per-capita administrative cost burden on remaining participants, and generates an ambiguous welfare effect whose sign depends on the relative size of the M group. The condition for a Pareto improvement in this case is derived at equation (B.10). Section B.3 draws out the policy implications across both cases. Throughout, welfare is evaluated using the expected utility of each type, with the social planner treating the equilibrium premium as the relevant price signal.

B.1 Welfare Effects under Adverse Selection

Consider the separating equilibrium illustrated in Figure 2. Baseline welfare is:

$$W_0 = \alpha V_L(S'_L) + (1-\alpha)\, V_H(S_H) \tag{B.1}$$

where α is the proportion of low-risk types, S'_L denotes the partial-coverage contract for L-types with coverage $\lambda_L < 1$, and S_H is the full-coverage contract for H-types.

Source of inefficiency. The separating equilibrium generates deadweight loss because L-types are under-insured relative to their first-best allocation. Following Einav et al. (2010) and Einav and Finkelstein (2011), who show how price variation identifies welfare losses in insurance markets, the welfare loss per L-type is:

$$DWL_L = V_L(S_L) - V_L(S'_L) \tag{B.2}$$

where S_L denotes full insurance at the actuarially fair premium $\Pi_L = \theta_L \bar{C}$.

Intervention: mandated minimum coverage. Consider a policy that mandates minimum coverage $\lambda_{\min} > \lambda_L$, potentially inducing pooling if separation is no longer feasible. For a pooling contract P with $\lambda_P \geq \lambda_{\min}$ to be Pareto improving, three conditions must hold simultaneously. First, the pooling premium must satisfy the zero-profit condition:

$$\Pi_P = \big[\alpha\theta_L + (1-\alpha)\theta_H\big]\lambda_P \bar{C} \tag{B.3}$$

Second, L-types must be weakly better off under pooling:

$$V_L(P) \;\geq\; V_L(S'_L) \tag{B.4}$$

Third, H-types must be compensable from the efficiency gains generated:

$$V_H(P) + T \;\geq\; V_H(S_H) \tag{B.5}$$

where T is a transfer funded by the welfare improvement. Combining these three conditions, the intervention generates a Pareto improvement if and only if the proportion of low-risk types exceeds the critical threshold:

$$\alpha \; > \; \alpha^* \; = \; \frac{\theta_H - \theta_L}{\theta_H - \theta_L + \dfrac{DWL_L}{\text{subsidy per H-type}}} \tag{B.6}$$

where the subsidy per H-type equals $(\Pi_P - \theta_H \lambda_P \bar{C})$ under pooling. When $\alpha > \alpha^*$, the efficiency gains from reducing under-insurance among L-types are large enough to finance the cross-subsidy required to make H-types no worse off.

B.2 Welfare Effects under the Double-Margin Mechanism

We analyse welfare under two distinct policy experiments. The first is an expansion of public coverage G holding the eligibility threshold Y_b fixed. The second is a reduction in the eligibility threshold to admit M types to public coverage. These experiments have different welfare properties and must be treated separately.

The population consists of L types with mass α, and H types with mass $(1-\alpha)$, so that $\alpha + (1-\alpha) = 1$. Following the introduction of M types in Section 4.1, the H population is subdivided: a fraction β are M types, with mass $(1-\alpha)\beta$, and the remaining fraction $(1-\beta)$ are Hl types, with mass $(1-\alpha)(1-\beta)$. All population shares therefore sum to one: $\alpha + (1-\alpha)(1-\beta) + (1-\alpha)\beta = 1$.

B.2.1 Welfare Effects under an Expansion of G with Eligibility Fixed

The initial equilibrium features three groups. L types (mass α) purchase private insurance at the pooling contract P^{**}. Hl types (mass $(1-\alpha)(1-\beta)$) are at the margin of indifference between private and public coverage, with $V_{Hl}(P^{**}) = V_{Hl}(G)$, and a fraction ε' of them have already taken public coverage. M types (mass $(1-\alpha)\beta$) have entered the private market at P^{**} following crowd-in and are ineligible for public coverage since $Y_M > \bar{Y}_b$. Initial welfare, abstracting from the tax cost of G, is:

$$W_0 = \alpha V_L(P^{**}) + (1-\alpha)(1-\beta)\big[(1-\varepsilon')V_{Hl}(P^{**}) + \varepsilon' V_{Hl}(G)\big] + (1-\alpha)\beta\, V_M(P^{**}). \tag{B.1}$$

Now let G rise, holding Y_b fixed. By the feedback chain (d) of Definition 1, ε' rises, more Hl types exit the private pool, and the zero-profit condition drives the pooling premium down. Since $\theta_M = \theta_H$, the private pool at the new equilibrium P^{***} satisfies:

$$\Pi_{P^{***}} = k + \frac{\alpha\theta_L + (1-\alpha)\big[(1-\beta)(1-\varepsilon'') + \beta\big]\theta_H}{\alpha + (1-\alpha)\big[(1-\beta)(1-\varepsilon'') + \beta\big]}\,\lambda_{P^{***}}\bar{C}(q_h), \tag{B.2}$$

where $\varepsilon'' > \varepsilon'$. The mass of M types in the pool is $(1-\alpha)\beta$ before and after the expansion, since M types remain ineligible for public coverage. The zero-profit condition (B.2) is strictly decreasing in ε'' since $\theta_H > \theta_L$: the premium falls as more Hl types exit.

All groups remaining in the private market are therefore better off. The welfare change has three components. L types gain from the lower premium:

$$\Delta W_L = \alpha\big[V_L(P^{***}) - V_L(P^{**})\big] > 0. \tag{B.3}$$

For Hl types, those moving to public coverage at the higher G gain by construction, and inframarginal Hl types remaining in the private pool gain from improved contract terms. Since $V_{Hl}(G) = V_{Hl}(P^{**})$ at the margin of indifference, the second term below vanishes:

$$\Delta W_{Hl} = (1-\alpha)(1-\beta)\big[(1-\varepsilon'')\big(V_{Hl}(P^{***}) - V_{Hl}(P^{**})\big) + (\varepsilon''-\varepsilon')\big(V_{Hl}(G) - V_{Hl}(P^{**})\big)\big] \geq 0. \tag{B.4}$$

M types gain from the lower premium:

$$\Delta W_M = (1-\alpha)\beta\big[V_M(P^{***}) - V_M(P^{**})\big] > 0. \tag{B.5}$$

The total welfare change is $\Delta W = \Delta W_L + \Delta W_{Hl} + \Delta W_M \geq 0$. Abstracting from the tax cost of financing the higher G, every term is non-negative and the welfare gain is unambiguously Pareto improving. Once the tax cost of G is accounted for, the net welfare gain is positive if and only if the efficiency gains from improved private contract terms exceed the additional fiscal cost — a standard public finance condition that does not depend on the selection structure.

B.2.2 Welfare Effects under a Reduction in the Eligibility Threshold to Admit M Types to Public Coverage

Now suppose the eligibility threshold is lowered so that $Y_M < Y_b$, making M types eligible for public coverage at G. This experiment has a qualitatively different welfare structure from Experiment 1. M types exit the private pool discontinuously, moving directly to public coverage at G. Since $\theta_M = \theta_H > \theta_L$, removing M types from the pool lowers the average risk weight and therefore reduces the pooling premium: L types and remaining Hl types in the private market are better off on the private market margin. The welfare cost of the experiment is instead fiscal: extending public eligibility to M types requires the public system to cover their actuarial risk $\theta_H G$, and this cost must be financed through taxation falling on all groups.

With M types on public coverage the private pool contracts to L types and the remaining fraction $(1-\varepsilon'')$ of Hl types. The zero-profit condition becomes:

$$\Pi_{\text{new}} = k + \frac{\alpha\theta_L + (1-\alpha)(1-\beta)(1-\varepsilon'')\theta_H}{\alpha + (1-\alpha)(1-\beta)(1-\varepsilon'')}\,\lambda_{\text{new}}\bar{C}(q_h), \tag{B.6}$$

where ε'' adjusts to maintain the Hl-type indifference condition $V_{Hl}(P_{\text{new}}) = V_{Hl}(G)$ at the new premium. Comparing (B.6) with (B.2), the removal of M types eliminates mass $(1-\alpha)\beta$ carrying risk θ_H from both numerator and denominator of the risk weight. Since $\theta_H > \theta_L$, removing these higher-risk types lowers the average risk weight in the pool, so $\Pi_{\text{new}} < \Pi_{P^{**}}$: the private premium falls. The administrative cost term k enters identically on both sides and cancels in this comparison, playing no role in the premium movement.

The welfare change has three components, abstracting from the tax cost of financing the extended public eligibility. M types gain from accessing public coverage:

$$\Delta W_M = (1-\alpha)\beta\big[V_M(G) - V_M(P^{**})\big] > 0. \tag{B.7}$$

L types benefit from the lower private premium:

$$\Delta W_L = \alpha\big[V_L(\Pi_{\text{new}}) - V_L(P^{**})\big] > 0. \tag{B.8}$$

Hl types remaining in the private pool also benefit from the lower premium, while those who move to public coverage are indifferent at the margin by the indifference condition:

$$\Delta W_{Hl} = (1-\alpha)(1-\beta)(1-\varepsilon'')\big[V_{Hl}(\Pi_{\text{new}}) - V_{Hl}(P^{**})\big] \geq 0. \tag{B.9}$$

Abstracting from the tax cost of financing G, all three terms are non-negative. The relevant welfare question is therefore whether the gains to all groups from the private premium reduction and from M types accessing public coverage exceed the fiscal cost of extending public eligibility. Letting $T = (1-\alpha)\beta\,\theta_H G$ denote the actuarial cost of covering M types, the condition for a welfare improvement is:

$$\begin{aligned}(1-\alpha)\beta\big[V_M(G) - V_M(P^{**})\big] &+ \alpha\big[V_L(\Pi_{\text{new}}) - V_L(P^{**})\big] \\ &+ (1-\alpha)(1-\beta)(1-\varepsilon'')\big[V_{Hl}(\Pi_{\text{new}}) - V_{Hl}(P^{**})\big] > (1-\alpha)\beta\,\theta_H G. \end{aligned} \tag{B.10}$$

Since all terms on the left-hand side are positive gains, condition (B.10) is satisfied whenever the combined welfare gains to all three groups exceed the actuarial cost of extending public coverage to M types.

Corollary 1. *Condition (B.10) is more likely to hold when β is small. When M types constitute a small fraction of the H population, the fiscal cost $(1-\alpha)\beta\,\theta_H G$ of extending public eligibility is modest relative to the private premium gains accruing to L and Hl types from the improved risk pool. The optimal scope of a public eligibility expansion therefore depends on the net gain to M types from public coverage relative to their actuarial cost, and on the empirical distribution of β, income, and risk aversion across the population.*

B.3 Policy Implications

The welfare analysis yields two main policy implications. In markets with adverse selection, condition (B.6) shows that mandated minimum coverage can generate a Pareto improvement when the proportion of low-risk types is sufficiently large to sustain the cross-subsidy required to compensate H-types. In markets where the double-margin mechanism is active, condition (B.10) shows that expanding public eligibility to cover M-types can improve welfare when the social surplus from avoiding administrative loading exceeds the premium deterioration imposed on L-types — a condition more likely to hold when the M group is small relative to the privately insured population. In both cases the relevant parameters are empirical quantities, and the design of public insurance expansions should be conditioned on evidence about the composition of the risk pool and the distribution of income relative to the eligibility threshold.

Under NHS-type universal coverage the welfare calculus is fundamentally different: since there are no uninsured individuals, the extensive-margin welfare gain from condition (B.10) does not arise, and the relevant trade-off is between the fiscal cost of improved public provision and the equity implications of differential access to supplementary private coverage, as discussed in Section 4.2.

References

[1] Alessie, R.J.M., Angelini, V., Mierau, J.O., and Viluma, L. (2020). Moral hazard and selection for voluntary deductibles. *Health Economics*, 29(10), 1251–1269.

[2] Boone, J., and Schottmüller, C. (2017). Health insurance without single crossing: Why healthy people have high coverage. *The Economic Journal*, 127(599), 84–105.

[3] Card, D. and Shore-Sheppard, L. (2004) Using discontinous eligibitlity rules to identify the effects on the federal medicaid expansions on low-income children, *Review of Economics and Statistics*, 86(3), 752-766

[4] Chade, H., Marone, V.R., Starc, A.,Swinkels, J., (2022), Mutidimensional screening and menu design in health insurance markets, NBER Working Paper 30542, NBER, Cambridge MA http://www.nber.org/papers/w30542

[5] Chetty, R. and Saez, E. (2017) Optimal Taxation and Social Insurance with Endogenous Private Insurance *American Economic Journal: Economic Policy*. 2(2): 85–114

[6] Cutler, D., and J. Gruber. 1996. "Does Public Insurance Crowd Out Private Insurance?" *Quarterly Journal of Economics* 111: 391–430.

[7] deMeza, D., and Webb, D.C. (2001). Advantageous selection in insurance markets. *RAND Journal of Economics*, 32(2), 249–262.

[8] Dubay, L.C., Kenney, G., 1996. Revisiting the issues: the effects of Medicaid expansions on insurance coverage of children. *The Future of Children*, 6 (1), 152–161.

[9] Dubay, L.C., Kenney, G., 1997. Did Medicaid expansions for pregnant women crowd-out private insurance? *Health Affairs*, 16 (1), 185–193.

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

[11] Einav, L., and Finkelstein, A. (2011). Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives*, 25(1), 115–138.

[12] Encinosa, W.E. (2003). The economics of public and private health insurance: Crowd-out, crowd-in, and push-out effects. *Portuguese Economic Journal*, 2(2), 71–86.

[13] Fang, H., Keane, M.P., and Silverman, D. (2008). Sources of advantageous selection: Evidence from the Medigap insurance market. *Journal of Political Economy*, 116(2), 303–350.

[14] Geruso, M., and Layton, T.J. (2017). Selection in health insurance markets and its policy remedies. *Journal of Economic Perspectives*, 31(4), 23–50.

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

[16] Golosov, M. and Tsyvinski, A, (2015) Policy implications of dynamic public finance. *Annual Review of Economics*, 7, 147-171

[17] Gruber, J. and Simon, K., (2008), Crowd-out 10 years later: Have recent public insurance expansions crowded out private insurance?. *Journal of Health Economics*, 27, 201-217

[18] Handel, B., Kolstad, J., Minten, T., and Spinnewijn, J. (2024). The socioeconomic distribution of choice quality: Evidence from health insurance in the Netherlands. *American Economic Review: Insights*, 6(3), 395–412.

[19] Olivella, P., and Vera-Hernández, M. (2013). Testing for asymmetric information in private health insurance. *Economic Journal*, 123(567), 96–130.

[20] Remmerswaal, M., Boone, J., and Douven, R. (2023). Minimum generosity levels in a competitive health insurance market. *Journal of Health Economics*, 90, 102791.

[21] Rothschild, M., and Stiglitz, J.E. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4), 629–649.

[22] Shore-Sheppard, L., Buchmueller, T., Jensen, G., (2000). Medicaid and crowding out of private insurance: a re-examination using firm-level data. *Journal of Health Economics*, 19 (1), 61–91.

[23] Thorpe, K., Florence, C., (1998). Health insurance coverage among children: the role of expanded Medicaid coverage. *Inquiry*, 35 (4), 369–379.

[24] Veiga, A. and Weyl, E.G., (2016) Product design in selection markets, *Quarterly Journal of Economics*, 131, 1007-1056

[25] Weyl, E.G. and Veiga, A., (2017) Pricing institutions and the welfare cost of adverse selection *American Economic Journal: Microeconomics*, 9(2), 139-148

[26] van Kleef, R.C., van Vliet, R.C.J.A., and van de Ven, W.P.M.M. (2020). Risk equalization in competitive health insurance markets: Theory and practice in the Netherlands. *Health Economics*, 29(1), 1–14.