Data Sharing with Dynamic Contracts*

Alessandro Dovis
University of Pennsylvania
and NBER
adovis@upenn.edu

Paolo Martellini
NYU Stern
and NBER
paolo.martellini@nyu.edu

August 2026

Abstract

This paper studies optimal information disclosure in dynamic economies with income risk, where an incumbent firm learns about the consumer's persistent type and its competitors do not. When the incumbent can commit to long-term contracts, optimal disclosure reveals no information. Without commitment, no cross-subsidization is feasible under any disclosure policy due to adverse selection. Partial disclosure is typically optimal, and it implements intertemporal consumption smoothing. We apply the model to credit markets, studying how information sharing should respond to fintech entry, and to health insurance, asking how many firms should have access to a policyholder's health record.

First version: March 2024. This paper previously circulated under the title "Long-Term Contracts, Commitment, and Optimal Information Disclosure". We thank Nageeb Ali, Matteo Camboni, Marco Bassetto, V.V. Chari, Hal Cole, Piero Gottardi, Natalia Kovrijnykh, Alessandro Lizzeri, Humberto Moreira, Juan Pablo Nicolini, Alessandro Pavan, Jacopo Perego, Marzena Rostek, Maher Said, Venky Venkateswaran, Alexander von Hafften, and Ariel Zetlin-Jones for useful comments. Mariia Elkina provided superb research assistance.

1 Introduction

In many markets, firms engage in long-term relationships with consumers. By observing the history of transactions, incumbent firms acquire information that competitors do not have access to: health and car insurers gradually learn their customers' health or driving records; credit card companies infer repayment probability from spending patterns; employers know more about their employees than the public record shows.[1] This ex-post informational monopoly gives incumbents an advantage over the competition.

This paper studies to what extent incumbent firms should be required to share what they know. A growing number of policies regulate the disclosure of proprietary data. A prominent example is the European Union's General Data Protection Regulation (GDPR), which grants consumers control over the data collected by firms about them; in banking, the same principle underlies *open banking* initiatives, which require banks to make customers' financial histories available to competitors upon request.[2] We characterize the disclosure policy that maximizes ex-ante welfare, highlighting how it shapes competition and consumption insurance.

We address this question in a two-period economy in which firms compete to attract consumers and the incumbent learns the consumer's type over time. Consumers are risk averse, value insurance across states and over time, and can always switch firms in the second period. When firms can commit to the terms of the contract, the optimal disclosure policy reveals no information about consumers' histories, to minimize the value of the outside option for high-income consumers ("high types") and maximize the value obtained by low-income consumers ("low types"). When firms cannot commit, cross-subsidization across types is infeasible under any public disclosure policy. Partial disclosure can improve welfare by facilitating intertemporal consumption smoothing between the first and second period for high types. We apply the model to two current policy debates: open banking in credit markets and restrictions on the use of histories in health insurance. In the first context, we find that the arrival of more attractive entrants (e.g. fintech) flips the relationship between pool quality and optimal information disclosure. In the second, we establish conditions under which it is optimal for one firm or no firm at

[1] For example, Kahn (2013) finds supporting evidence for the labor market, Cohen (2012) for the auto insurance market, and Ioannidou and Ongena (2010) for the corporate loan market.

[2] In the United States, legislation inspired by the GDPR has been adopted in several states, such as California's Consumer Privacy Act (2018), although its scope in the banking sector remains limited. At the federal level, the Personal Financial Data Rights Rule finalized by the Consumer Financial Protection Bureau (CFPB) in 2024—currently facing legal challenges—seeks to extend open banking principles to U.S. financial institutions. Similar initiatives have been implemented in other jurisdictions, including the United Kingdom, Australia, and Brazil.

all to condition the contracts offered on the consumer's health records.

We study a two-period economy where risk-averse consumers trade with risk-neutral firms. Income takes two values, high or low, and is Markovian. In the first period, all agents share the same information and firms compete for consumers by offering contracts. At the end of the first period, the consumer and the incumbent learn the income realization. Competing firms (outsiders) observe only a public signal from the *disclosure policy*. In the second period, the incumbent can therefore condition its offer on the income realization, while outsiders' menus must satisfy incentive constraints, and consumers decide whether to switch firms. To ensure equilibrium existence (in pure strategies), we let firms withdraw posted contracts at a small cost, as in Netzer and Scheuer (2014).

A public disclosure policy maps the first-period income realization—our proxy for a consumer's type—into a signal observed by everyone. By controlling the composition of the pool of consumers with a given signal, the policy determines the most outsiders are willing to offer to high types in the second period; this outside option in turn shapes the contract the incumbent offers in the first period.

When firms can commit to long-term contracts, no disclosure is optimal. Commitment lets the incumbent deliver as much consumption smoothing and cross-subsidization as permitted by the high types' ex-post participation constraint, which tightens when information is disclosed (Hirshleifer (1971)). Despite the private information, the equilibrium consumption profile features the same pattern as in Harris and Holmstrom (1982) and Thomas and Worrall (1988): consumption is downward rigid and *back-loaded*.

Next, we characterize the equilibrium without firm commitment across periods. The incumbent can deliver any value to high types more cheaply than outsiders. This is because outsiders cannot distinguish high types from low types and must therefore either distort high-type consumption to screen out low types or subsidize low types. The incumbent can therefore always outbid outsiders for the high types, retaining them at their best outside offer and earning profits that shrink as the signal becomes more informative. Outsiders anticipate that any consumer willing to switch must be a low type. Without commitment, low types thus receive no transfers under any disclosure policy, in contrast with Netzer and Scheuer (2014), where all firms are equally uninformed and low types may be cross-subsidized. Ex-ante competition rebates the ex-post profits to consumers in the first period.

We then characterize the optimal public disclosure policy in such environment. Since low types' second-period consumption is independent of the information structure, the design of disclosure is driven entirely by intertemporal smoothing between the first- for second-period consumption for the high types. The optimal policy is a "*bad news*" system:

all high types receive a good signal, and only some low types do. The more low types are pooled with high types, the lower the high types' outside option and the higher the profits rebated in the first period. At the optimum, the equilibrium contract inverts the Harris and Holmstrom (1982) profile: consumption is upward rigid and *front-loaded*.

To set the stage for our applications, we characterize optimal disclosure in a static benchmark: an economy of spot contracts in which firms are symmetrically uninformed. Absent an informed incumbent, pooling contracts do cross-subsidize low types, but there are no ex-post rents to rebate, hence no intertemporal smoothing. When the share of high-income consumers is high, no disclosure is optimal: firms pool all consumers at common terms, subsidizing the low types. When the share is low, firms would screen consumers absent additional information, and some disclosure is needed to trigger cross-subsidization of the few low types that are assigned a good signal.

In our first application, we specialize the environment to bank lending: the incumbent bank observes its borrowers' repayment records, while competitors see only the credit information it is required to share. We derive two implications for open banking. First, lenders should share a coarse credit rating rather than full credit histories: the top rating pools all creditworthy borrowers together with some risky ones. A corollary is that data ownership should not be transferred to borrowers, contrary to what the GDPR's right to data portability and the related open banking reforms dictate.[3]

We also study how disclosure policy should change with the entry of fintech. We interpret fintech as outsiders that, in the second period, can offer additional benefits—such as faster origination or broader access—that give creditworthy borrowers a reason to switch. How information sharing should respond depends on the composition of the borrower pool. When most borrowers are creditworthy, information sharing should decrease. Absent better outsiders, consumption-smoothing considerations imply that some information should be shared: without it, loan terms would deteriorate too steeply within the relationship. With fintech, however, providing no information is optimal, because pooling the many creditworthy switchers with risky borrowers sustains cross-subsidization. When the pool is poor, by contrast, no information should be shared absent fintech, while information sharing should increase with fintech: only a finer rating induces entrants to cross-subsidize risky borrowers.

In our second application, we map the model to health insurance, where regulation often bans the use of experience rating, the incumbent's informational advantage, in fa-

[3]In principle, the GDPR's right to erasure could sever the incumbent's informational advantage altogether. In banking, erasure is impractical: lenders must retain records to comply with anti–money laundering and prudential regulations.

vor of community rating, implementing our static benchmark. We ask how many insurers should be able to act on a policyholder's health history. Experience rating exposes policyholders whose health deteriorates to steeper premiums, while community rating exacerbates adverse selection. The answer once again depends on the pool quality: when it is healthy, the cross-subsidization allowed by community rating dominates; when it is unhealthy, policyholders are better off with an informed incumbent, whose rents finance premium discounts early in the relationship. In such case, consumers want one and only one firm to know their type.

Related literature This paper contributes to the literature on insurance contracts under lack of commitment, in the spirit of the seminal contributions by Harris and Holmstrom (1982) and Thomas and Worrall (1988). This class of models has been used to study wage-tenure profiles, by Burdett and Coles (2003) and Balke and Lamadon (2022), and life-insurance contracts, by Hendel and Lizzeri (2003). In these models, the agent's outside option determines how much insurance is attainable. We contribute by assuming that the incumbent firm acquires more information over time about the consumer's persistent type than its competitors do, and we show how this shapes the consumer's outside option.[4] We then study how public disclosure policies affect potential competition and thereby the amount of insurance and consumption smoothing attainable in equilibrium.

In solving for the optimal disclosure policy, our paper connects to the literature on information design, as in Kamenica and Gentzkow (2011) and Bergemann et al. (2015). Closest to our work, Garcia and Tsur (2021) and Immorlica et al. (2022) study optimal information disclosure in adverse-selection economies, while Farinha Luz et al. (2023), Elliott et al. (2024), and Fainmesser et al. (2023) study how disclosure of consumer information shapes market equilibria. In all these papers, consumers and firms are short-lived, and firms are symmetrically uninformed.

The closest papers to ours are de Garidel-Thoron (2005) and Mukherjee (2008), which also study information disclosure in dynamic insurance economies where the incumbent gains an informational advantage. de Garidel-Thoron (2005)'s main analysis is under commitment on the firm's side, and he compares two polar cases—full and no disclosure—finding that no disclosure yields higher welfare. Our key contribution is to show that

[4] Sharpe (1990) studies a model with asymmetric learning between firms and long-term lending relationships. The equilibrium dynamics—for a given disclosure policy—are similar to those in our economy, but he does not study the optimal information disclosure. More recently, Kaniel and Orlov (2025) study a labor market in which the worker's quality gradually becomes public information and firms can extract rents by affecting the evolution of those beliefs via firing and retention decisions. Another strand of the literature allows for asymmetric information to arise between the agent and all the firms, e.g., Green (1987), or about the outside offer that the agent receives, e.g., Hopenhayn and Werning (2008).

some information disclosure is optimal when firms lack commitment. We further study when consumers would prefer no firm to have access to information, and relate these insights to recent policy debates in the health insurance and lending markets. Mukherjee (2008) studies an environment where second-period outsiders are more efficient than the incumbent, as in Section 8, and finds conditions under which full disclosure is optimal. Critical to his result are firm commitment and the incumbent's ability to penalize agents who leave. Absent either, information affects not only the resources the incumbent extracts but also how they are distributed across time and states, and full disclosure is never optimal.

Our paper also relates to studies of competition in frictional markets with adverse selection, such as Guerrieri (2008), Guerrieri et al. (2010), and Lester et al. (2019), where information is symmetric across firms. In our economy, ex-post competition is instead shaped by the incumbent's informational advantage.

Inspired by a growing debate, a recent literature studies open banking policies, mandates to share information with the objective to stimulate competition from fintech entrants; see Babina et al. (2024) and He et al. (2023). Kurlat and Scheuer (2021) and Farboodi et al. (2025) study environments in which lenders differ in their screening expertise; the latter also analyzes how data-sharing policies affect lender entry and credit-market outcomes. All of these papers restrict attention to static contracts, abstracting from intertemporal smoothing and its interaction with cross-subsidization. In Kovrijnykh et al. (2024), information sharing operates through a different channel: access to a borrower's credit history lets lenders infer the signals received by competitors, which can raise credit limits. There the entire history of accepted contracts is public, and the design of disclosure is not analyzed.

Our perspective on consumer data ownership differs from that in Jones and Tonetti (2020), where assigning data property rights to consumers improves the allocation by letting multiple firms use the same data—in line with the GDPR's *Right to data portability*—in Ali et al. (2023), where the volume of trade increases from consumers' voluntary disclosure of verifiable information about their valuations, and in Eeckhout and Veldkamp (2022), where an imbalance of data between competing firms generates markups that hurt consumers. We find a complementary channel: with insurance considerations, assigning full data ownership to consumers is welfare-reducing because it kills insurance within and across periods. Granting consumers the GDPR's *Right to be forgotten* can instead be beneficial, by removing the incumbent's informational advantage, but only if consumers can commit to data erasure upon signing a contract. This commitment is precisely what regulation implicitly provides in health insurance, where community rating bars insurers

from using policyholders' health histories. Handel et al. (2015) study when such bans are optimal in a static environment with symmetrically informed insurers. We extend their comparison to a dynamic one and find conditions under which granting the use to price consumers' histories to a single (incumbent) firm dominates both full disclosure and full data erasure.

The rest of the paper is organized as follows. Section 2 presents the environment. Section 3 shows how public disclosure affects potential competition and the second-period equilibrium. Sections 4 and 5 study optimal disclosure with and without firm commitment. Section 6 discusses extensions. Section 7 presents the static benchmark with symmetrically uninformed firms. Sections 8 and 9 develop the applications. Section 10 concludes. Omitted proofs are in Appendix B.

2 Environment

Consider an economy that lasts for two periods, $t = 1, 2$. There are two types of agents: consumers and firms. Each period is divided into two subperiods. In the first subperiod, consumers receive no income–this is simply a normalization. Consumers are uncertain about their income in the second subperiod, y, which can take a high value, y_H, or a low value, y_L. We assume that $y_1 \sim \pi_1(y_1)$ and $\pi_2(y_2|y_1)$ is the conditional probability of y_2 in the second period after a realization of y_1 in the first period. Thus, period 1 income is useful to forecast period 2 income. We denote consumers with a realization of y_H in period 1 as the high-income consumers or, simply, high types. We define Y_{2H} and Y_{2L} to be the expected income in period 2 conditional on having a high or low income realization in period 1, respectively,

$$Y_{2H} \equiv y_L + \pi_2(y_H|y_H)(y_H - y_L) > Y_{2L} \equiv y_L + \pi_2(y_H|y_L)(y_H - y_L).$$

We assume that $Y_{2H} > Y_{2L}$, so a higher income realization in period 1 forecasts a higher average income in period 2. We also define the (unconditional) expected income in period 1, $Y_1 = y_L + \pi_1(y_H)(y_H - y_L)$, and in period 2, $Y_2 = \sum_{i=L,H} \pi_1(y_i) Y_{2i}$. We assume that $Y_{2L} \leqslant Y_1 \leqslant Y_{2H}$, or, equivalently, that $\pi_2(y_H|y_L) \leqslant \pi_1(y_H) \leqslant \pi_2(y_H|y_H)$. Last, we assume that $\pi_2(y_H|y_1) \in (0, 1)$ for all y_1 so that there is residual uncertainty in period 2.

In the Appendix, we show how this restriction arises in an economy where consumers can be of one of two unobservable types, $\theta \in \{\theta_L, \theta_H\}$, that affect the probability distribution of income $p(y|\theta)$, with $p(y_H|\theta_H) > p(y_H|\theta_L)$, and agents learn about the consumer's type by observing income realizations. Such an economy can be exactly mapped into our

economy with $Y_1 = Y_2$, so the expected income in period 1 and the unconditional expected income in period 2 are the same. We choose to illustrate our results assuming income in the second subperiod follows a Markov chain to minimize the required notation.

Consumers are risk-averse and have common preferences over expected per-period consumption, C, given by $u(C)$, with u strictly increasing and strictly concave. Consumption within the period is aggregated linearly: consumers value consumption in the first subperiod, c_{am}, at factor 1 and in the second subperiod at factor $\beta < 1$. We assume that consumers have limited liability so that their consumption must be non-negative in each subperiod. As we show below, it is never optimal to offer strictly positive consumption in the second subperiod after the realization of y_L. We can then denote consumption after y_H simply as c_{pm} and express, $C_1 = c_{am,1} + \beta\pi_1(y_H)\, c_{pm,1}$ in period 1 and $C_2 = c_{am,2} + \beta\pi_2(y_H|y_1)\, c_{pm,2}$ in period 2.

Firms are risk-neutral and value income in both subperiods at factor 1. To simplify notation, without loss of generality, we assume no discounting between period 1 and period 2 for either firms or consumers.

We briefly comment on the main features of our preferences. The differential discounting within each period generates gains from trade: within each period, firms aim at front-loading their transfers to consumers, subject to the uncertainty about income realizations. If within-period discounting was the same for consumers and firms, perfect separation would be costless, effectively eliminating the information advantage of the incumbent. Risk-aversion generates a utility loss whenever expected consumption is unequally distributed across time and state, period 1 and period 2 after each realization of y_1. Absent risk-aversion, consumers only care about the present value of consumption which is always equal to the present value of income due to competition ex-ante. Thus, both differential discounting and risk-aversion are needed for information design to affect consumer welfare.

Our environment captures the same forces as a standard Rothschild-Stiglitz economy within each period: the usual trade-off between separating and pooling consumer types emerges. Instead of assuming insurance motives *within a period*, we adopt a linear aggregation, augmented with differential discounting between consumers and firms. Just as in standard insurance problems, in our model the marginal utility of resources varies across states (here, subperiods) for consumers, but not for firms. We can then interpret the first subperiod as the high marginal utility state, and the second subperiod after a high-income realization as the low marginal utility state. A key benefit of our formulation is that it allows us to obtain closed-form expressions for the allocation and a sharper char-

acterization of the separating and pooling regions.[5] Last, concavity of utility *across periods* generates insurance motives against variations in consumption induced by the income realization in the first period. (A previous version of the paper, Dovis and Martellini (2024), considers a more canonical two-period insurance economy that delivers all the same results, except for Lemma 3 which we verify in numerical examples and for which we provide a sufficient condition.)

Information and market structure Firms compete to attract consumers by offering contracts that specify consumption levels conditional on all the available information at the time. At the beginning of period 1, all agents share the same information. Firms simultaneously offer contracts $(c_{1,am}, c_{1,pm})$ to consumers, who choose one firm to sign with. We will call the firm chosen by the consumer in period 1 the *incumbent*. The firm pays consumption $c_{1,am}$. At the end of period 1, income y_1 is realized and it is observed by the consumer and the incumbent. The payment $c_{1,pm}$ is made. The *outsiders* (i.e., the other firms) do not observe y_1 directly but observe a public signal m in some signal space M. Public signals are distributed according to some distribution $\mu \in \Delta(M)$, where $\mu(m|y_1)$ denotes the share of consumers with income y_1 that receive signal m. We will refer to (M, μ) as the *public disclosure policy*.

At the beginning of period 2, the incumbent acts as a Stackelberg leader and offers a contract $(c_{2,am}, c_{2,pm})$ conditional on the consumer's history (y_1, m). Outsiders observe the distribution of contracts offered by the incumbent and can poach consumers by offering a menu of contracts $(c_{2,am}, c_{2,pm})$ to all consumers with a given public signal m. Critically, outsiders do not observe the realization of y_1 and the contract offered by the incumbent to a particular consumer. All firms observe the contracts offered and decide whether to withdraw their contracts at a cost $\varepsilon \geqslant 0$. Finally, consumers choose the contract that maximizes their utility among those that are left after the withdrawal stage. A consumer with history (y_1, m) can either accept the single contract offered by the incumbent or select one of the contracts in the menu offered by the outsiders to consumers with history m. If there are contracts that offer the same value, we consider the following tie-breaking rule: if the incumbent is among those offering the contract with the highest value, then the consumer stays with the incumbent; if not, the consumers are evenly split among the outsiders with the same offer.[6]

[5] In doing so, we connect with the literature on redistribution in linear environments (e.g., Dworczak et al. (2021)).

[6] The selection rule is without loss of generality, because the incumbent has an informational advantage and can always attract all consumers by offering a little more consumption and still making positive profits. Our selection rule effectively makes the incumbent's choice set closed.

We always assume that consumers can walk away from the incumbent and sign with another firm.[7] On the firm side, we consider two cases. First, we assume, as a benchmark, that the incumbent can commit in period 1 to the terms of the period 2 contract. Second, we assume that also the firm cannot commit ex-ante to the terms of the period 2 contract. Under both assumptions on commitment, we characterize the equilibrium outcome for a given public disclosure policy and then characterize the one that maximizes ex-ante welfare.

Remark. In setting up the interaction among firms in period 2, we follow Netzer and Scheuer (2014) and allow for a stage in which each firm observes other firms' contracts and can decide to withdraw its own at a small cost $\varepsilon > 0$. We highlight a similarity and a key distinction with respect to their model. As in Netzer and Scheuer (2014), the presence of a withdrawal stage guarantees that there exists a continuation equilibrium in pure strategies in period 2 for any signal m. In particular, the ability to withdraw the offered (menu of) contracts allows for the existence of equilibrium outcomes that feature cross-subsidization. The withdrawal stage rules out deviations in which competitors offer a "cream-skimming" contract to the high types. These deviations are at the root of the non-existence of pure strategy equilibria in Rothschild and Stiglitz (1976).

In Netzer and Scheuer (2014), the presence of a strictly positive withdrawal cost, $\varepsilon > 0$, guarantees uniqueness of continuation equilibrium outcomes by ruling out those that generate positive profits for some firms. Our model differs from theirs due to the presence of an informed incumbent. For us, if $\varepsilon > 0$, uniqueness of equilibrium outcomes–and, in fact, even equilibrium existence–additionally requires that the incumbent acts as a Stackelberg leader. This requirement is naturally satisfied when the incumbent can commit, given that the terms of the contract are predetermined at the beginning of period 2. We maintain the Stackelberg assumption also under lack of firm commitment.[8] We formalize these arguments in Appendix B.2 and B.5 for the cases with and without firm commitment.

[7] If both the incumbent and the consumer can commit to the terms of the contract in period 2 and to staying with the firm, respectively, then the public disclosure policy is irrelevant and $C_t(y^t) = c_{t,am}(y^t) = \frac{Y_1+Y_2}{2}$ and $c_{t,pm}(y^t) = 0$.

[8] An alternative approach would be to set $\varepsilon = 0$ without restrictions on whether the incumbent is a Stackelberg leader or all firms act simultaneously in period 2. The equilibrium outcome we characterize below would correspond to one of the many possible when $\varepsilon = 0$, specifically the one that minimizes period 2 profits for the incumbent.

3 Miyazaki-Wilson value

We start by characterizing the equilibrium value in period 2 that would arise absent an informed incumbent, following Netzer and Scheuer (2014). The value associated to this equilibrium outcome will serve as a critical outside option that will shape the contract offered by the informed incumbent.

Consider agents with a signal m, and denote by s the fraction of consumers that have $y_1 = y_H$:

$$s = \frac{\mu(m|y_H)\,\pi_1(y_H)}{\sum_{y_1}\mu(m|y_1)\,\pi_1(y_1)}. \tag{1}$$

To characterize the continuation equilibrium, it is useful to define the maximal expected consumption that can be offered by the outsider to a consumer with history (y_H, m). The outsider chooses a consumption profile for the high-income consumers, (c_{am}, c_{pm}) and expected consumption intended for consumers with history (y_L, m), the low-income consumers, to maximize:

$$C_H^o(s) = \max_{c_{am}, c_{pm}, C_L^o} c_{am} + \beta\pi_2(y_H|y_H)\,c_{pm} \tag{2}$$

subject to the outsider's nonnegative profit condition,

$$s\,[Y_{2H} - (c_{am} + \pi_2(y_H|y_H)\,c_{pm})] + (1-s)\,[Y_{2L} - C_L^o] \geqslant 0,$$

the incentive compatibility constraint for low-income consumers,

$$C_L^o \geqslant c_{am} + \beta\pi_2(y_H|y_L)\,c_{pm},$$

and the participation constraint for low-income consumers,

$$C_L^o \geqslant Y_{2L}. \tag{3}$$

We denote by $V_H^o = u(C_H^o)$ and $V_L^o = u(C_L^o)$ the associated consumers' values.[9] Throughout the paper, we say that consumption is distorted within the period whenever there is strictly positive consumption in the pm. In our description of (2) we impose two properties of the solution. First, the consumption of low types is undistorted. Second, consumption in the second subperiod after a realization of y_L is always set to 0, hence omitted

[9] To ease notation, in what follows we omit the dependence of C^o and V^o on s when the latter is kept fixed.

Figure 1: Expected consumption in Miyazaki-Wilson program

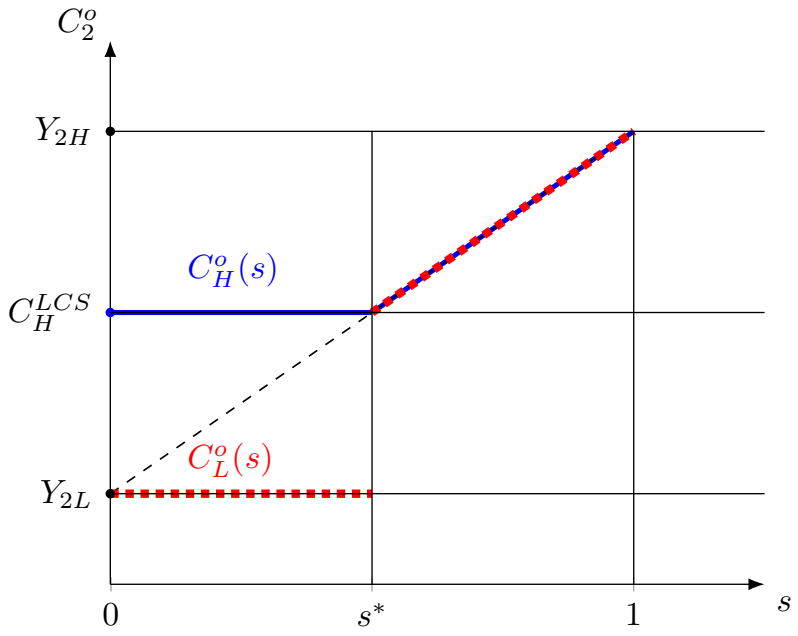

from the problem.[10]

The program (2) is what Netzer and Scheuer (2014) term the Miyazaki--Wilson program, after Miyazaki (1977) and Wilson (1977). In Netzer and Scheuer (2014), given the absence of an incumbent with information advantage, the solution to this program characterizes the unique equilibrium outcome of the game in period 2.

We now characterize the solution to the problem. Let

$$\kappa \equiv \frac{\pi_2(y_H|y_H) - \pi_2(y_H|y_L)}{\pi_2(y_H|y_H) - \pi_2(y_H|y_L)\beta}.$$

Lemma 1. *Let $s^* = \beta\kappa < 1$. If $s < s^*$, then the optimal allocation in (2) is the least-cost-separating contract:*

$$c_{am} = \pi_2(y_H|y_H)(y_H - y_L)(1-\kappa) + y_L, \quad c_{pm} = \kappa(y_H - y_L). \tag{4}$$

[10]The first simplification follows from the fact that the incentive compatibility constraint for the high type is never binding. To see why the second simplification is without loss, suppose consumption in the pm after y_L is strictly positive. Then, reducing it by an arbitrary small amount $\delta > 0$ and increasing c_{am} by $\beta\delta\pi_2(y_L|y_L)$ constitutes a profitable deviation, since it leaves the incentive compatibility constraints unaffected, it relaxes the nonnegative profit condition, and it increases the value of the objective function as $\beta\delta\pi_2(y_L|y_L) - \beta\delta\pi_2(y_L|y_H) > \beta\delta\pi_2(y_L|y_L) - \beta\delta\pi_2(y_L|y_L) = 0$.

and $C_L^o = Y_{2L}$. *If* $s \geqslant s^*$, *the optimal allocation in (2) features pooling:*

$$c_{am} = C_L^o = sY_{2H} + (1-s)\,Y_{2L}, \quad c_{pm} = 0. \tag{5}$$

If the share of high-income consumers is low, the participation constraint for the low-income consumers (3) is binding and the optimal solution is the *least-cost-separating* allocation with no cross-subsidization across types and expected consumption C_H^{LCS} (value $V_H^{LCS} = u\left(C_H^{LCS}\right)$) for the high-income type.[11] Since the participation constraint for low-income consumers binds, $C_L^o = Y_{2L}$, the allocation for the high-income consumers can be solved by combining incentive compatibility and zero-profits, obtaining (4). The allocation of high-income consumers is distorted. That is, they receive strictly positive consumption in the pm despite that being discounted at a factor $\beta < 1$. This distortion is necessary to prevent low-income consumers–who have a lower probability of having high-income in the pm–from pretending to be high-income types. This is equivalent to imperfect insurance offered to the low risk-type in a standard insurance economy.

If instead the share s is sufficiently high, the solution features cross-subsidization, $C_L^o > Y_{2L}$, as illustrated in Figure 1. This is because it is cheaper to provide a subsidy to the (few) low-income consumers than to distort the allocation for the high-income ones. In this case, the allocation satisfies perfect pooling across types. The allocation for high-income consumers is not distorted, but their expected consumption is lower than under full information unless $s = 1$.

We further characterize the minimal value of s that sustains pooling, $s^* = \beta\kappa$. Intuitively, s^* is increasing in β since more patient agents face a lower cost from delayed consumption, facilitating separation. Similarly, s^* is increasing in $\pi_2(y_H|y_H) - \pi_2(y_H|y_L)$: a larger spread between types relaxes the incentive compatibility constraint and lowers the adverse selection distortions of a separating contract.[12]

4 One-sided commitment

As a benchmark, we first consider the long-term equilibrium contract and the optimal public disclosure policy in the presence of an incumbent firm that can commit to the continuation contract, a setting closely related to existing results in the dynamic-contracts literature. We show that the optimal disclosure aims at minimizing the outside option for

[11] Formally, the least-cost-separating allocation solves a restricted version of (2), where the participation constraint (3) must hold with equality.

[12] Define $\Delta_\pi \equiv \pi_2(y_H|y_H) - \pi_2(y_H|y_L)$. We can then rewrite $s^* = \frac{\beta}{1+(1-\beta)\pi_2(y_H|y_L)/\Delta_\pi}$, from which both comparative statics are immediate.

the consumer that draws y_H in period 1 to maximize the degree of cross-subsidization between types. The disclosure policy that best relaxes the ex-post participation constraint for the high type reveals no information.

The potential competition of outsiders requires that the optimal contract satisfies the following ex-post participation constraints:

Lemma 2. *The optimal contract* $\{c_{1,am}, c_{1,pm}, c_{2,am}(y_1, m), c_{2,pm}(y_1, m)\}$ *satisfies*

$$c_{2,am}(y_H, m) + \beta\pi_2(y_H|y_H)\, c_{2,pm}(y_H, m) \geqslant C_H^o(s(m)), \tag{6}$$

$$c_{2,am}(y_L, m) + \beta\pi_2(y_H|y_L)\, c_{2,pm}(y_L, m) \geqslant Y_{2L}. \tag{7}$$

Moreover, the incumbent withdraws its offer if outsiders offer a cream-skimming contract intended for the high type.

It is clear that the equilibrium contract features undistorted consumption in period 1, $c_{1,am} = C_1$ and $c_{1,pm} = 0$, and in period 2 contingent on (y_1, m), i.e., $c_{2,pm}(y_1, m) = 0$ for all (y_1, m). Given a public disclosure policy (M, μ), the optimal contract offered in period 1 solves

$$\max_{C_1, C_2} u(C_1) + \sum_{y_1} \sum_{m} \mu(m|y_1)\, \pi_1(y_1)\, u(C_2(y_1, m)) \tag{8}$$

subject to the firm's nonnegative-profit condition,

$$Y_1 - C_1 + \sum_{i=L,H} \sum_{m} \mu(m|y_i)\, \pi_1(y_i)\, (Y_{2i} - C_2(y_i, m)) \geqslant 0,$$

and the ex-post participation constraints for high- and low-income consumers (6) and (7), respectively.

The participation constraint for the low-income consumers (7) is slack because the incumbent wants to subsidize consumption of the low types in period 2.

Optimal disclosure policy We now turn to the design of the optimal public disclosure policy. Before proceeding, we state the following Lemma that will be useful to characterize the optimum:

Lemma 3. $C_H^o(s)$ *is convex.*

That is, the incumbent's cost of providing utility $V_H^o(s)$ in period 2, $C_H^o(s)$, is convex in the share of high-income consumers. This follows immediately from Lemma 1 and it is illustrated in Figure 1. This lemma guarantees that the maximal profits the incumbent

firm makes on high-income consumers do not increase if those consumers are associated with more than one signal. Thus, information disclosure simply reduces to choosing the fraction of low-income consumers that receive the same signal as high-income consumers.

Proposition 1. *When the firm has commitment, the optimal disclosure policy reveals no information.*

The optimal public disclosure policy aims at minimizing the outside option for high-income consumers in period 2 to maximize the degree of cross-subsidization between types. Revealing no information, that is, assigning the same signal to all consumers, is optimal because it maximizes the distortions that the outsiders must impose on the high-income consumers to separate them from the low-income ones. These high distortions lower the value for the high-income consumers that can be offered by the outsiders and allow for greater cross-subsidization.

We can further characterize the optimal allocation under one-sided commitment. The optimal allocation has

$$C_1 = C_2(y_L) \leqslant C_2(y_H) = \max\{C_H^o, C_1\}, \tag{9}$$

where the inequality is strict if the ex-post participation constraint (6) is binding. As in Harris and Holmstrom (1982) and Thomas and Worrall (1988), the equilibrium allocation has perfect consumption smoothing between periods 1 and 2 after a low income realization. If (6) is binding, consumption must be increased after a high income realization to retain the high-income consumers, consistent with the evidence in Hendel and Lizzeri (2003) and Balke and Lamadon (2022) in the life insurance and labor market contexts, respectively.

5 Two-sided lack of commitment

We now turn to our main case: neither side can commit to contract terms beyond the current period. We first show that no disclosure policy can sustain cross-subsidization in period 2: low-types consume $C_2(y_L, m) = Y_{2L}$ regardless of (M, μ). Disclosure instead governs intertemporal smoothing for the high-income consumers: the incumbent's informational advantage determines the profits it extracts from them in period 2, and period-1 competition rebates those profits up front. The more information is disclosed, the lower the incumbent's ex-post profits and the more consumption tilts toward period 2.

The optimal disclosure policy can be described by a two-signal system, $M = \{g, b\}$. All high-income consumers receive the good signal g. Low-income consumers receive a good

signal with probability $1-\mu$ and a bad signal with probability μ. Thus, μ measures the degree of informativeness of the disclosure policy. If $\mu = 1$, there is perfect information, while if $\mu = 0$, the outsiders have no information in addition to their prior. We show that, perhaps surprisingly, the amount of information disclosed is increasing in the share of high-income consumers, $\pi_1(y_H)$.

Outcome in period 2

We characterize the outcome by backward induction starting from the terminal period 2. Agents have history y_1 and a signal $m \sim \mu(y_1)$. We model the incumbent as a Stackelberg leader that offers its contract to existing consumers, mimicking the timing in the case with one-sided commitment.

The next lemma characterizes the continuation equilibrium outcome for a given public disclosure policy:[13]

Lemma 4. *For any signal m, consumption is undistorted in period 2, that is, $c_{2,pm} = 0$. There is no cross-subsidization from the high-income consumers to the low-income consumers and the latter always consume $C_2(y_L, m) = Y_{2L}$. The consumption of the high-income consumers is*

$$C_2(y_H, m) = C_H^o(s(m)).$$

The incumbent must offer $C_H^o(s(m))$ to retain consumers with history (y_H, m). Unless the signal is perfectly informative, the incumbent profits from doing so since $C_H^o(s(m)) < Y_{2H}$. The reason is that the incumbent knows the period-1 realization, so its offer need not satisfy the incentive-compatibility constraint that forces outsiders to distort allocations or cross-subsidize. The incumbent therefore delivers the same value at strictly lower cost.

Low-income consumers always consume Y_{2L}, so firms make no profits on them. Even when a consumer with $y_1 = y_L$ receives the same good signal as her high-income peers, adverse selection implies that she still consumes Y_{2L}; her value is independent of the disclosure policy. Finally, except under perfect information revelation there is essentially no mobility in period 2: the incumbent strictly profits from retaining high-income consumers, and low-income consumers are indifferent between staying and switching. This explains why low-income consumers receive no transfers in equilibrium, regardless of their signal. Outsiders understand that any consumer willing to switch must be a low-income type, because the informed incumbent has a strict incentive to retain high-income

[13] We assume that the incumbent cannot discriminate among consumers with the same history. In Appendix D, we show how this assumption changes the specifics of the equilibrium allocation without affecting the main message. In particular, the optimal signal structure maintains the same features as in the restricted case.

consumers. This contrasts with the static economy with uniformly uninformed firms considered in Section 3, where firms are willing to offer pooling contracts that subsidize low-income consumers in order to reduce the cost of attracting high-income consumers.

Outcome in period 1

In period 1, firms compete for consumers: the equilibrium contract maximizes expected utility subject to a dynamic zero-profit condition. Anticipated period-2 profits on high-income consumers are competed away and rebated in period 1. The optimal contract has no static distortions—consumption is fully front-loaded, $C_1 = c_{1,am}$ and $c_{1,pm} = 0$—and solves

$$V_1 = \max_{C_1} u(C_1) + \pi_1(y_H) \sum_m \mu(m|y_H) V_H^o(s(m)) + \pi_1(y_L) u(Y_{2L}) \tag{10}$$

subject to

$$Y_1 - C_1 + \pi_1(y_H) \sum_m \mu(m|y_H) (Y_{2H} - C_H^o(s(m))) \geqslant 0.$$

It immediately follows that period 1 consumption is fully determined by the zero-profit condition holding with equality. We summarize the characterization of the equilibrium outcome given a public disclosure policy (M, μ) in the next lemma:

Lemma 5. *Given a public disclosure policy* (M, μ)*, the equilibrium outcome has*

$$C_1 = Y_1 + \pi_1(y_H) \sum_m \mu(m|y_H) [Y_{2H} - C_H^o(s(m))] \tag{11}$$

$$C_2(y_L, m) = Y_{2L} \tag{12}$$

$$C_2(y_H, m) = C_H^o(s(m)). \tag{13}$$

5.1 Optimal public disclosure policy

We now characterize the optimal public disclosure policy. Given Lemma 5, the choice of disclosure policy does not affect the consumption of the low-income consumer that always consumes Y_{2L} in period 2, but it can affect the consumption of high-income consumers in period 2 and consequently consumption in period 1. In fact, by affecting the share of high-income consumers with a given signal, the information designer can affect the outside options of these consumers in period 2 and therefore the retention value the incumbent offers them. Such values can range from the value of the least-cost-separating contract to $u(Y_{2H})$, because that is the range of V_H^o, as illustrated in Figure 1.

The optimal disclosure policy depends on whether the following condition is satisfied:

$$C_H^o(\pi_1(y_H)) > Y_1 + \pi_1(y_H)(Y_{2H} - C_H^o(\pi_1(y_H))). \tag{14}$$

That is, if under no information disclosure, expected consumption in period 2 after $y_1 = y_H$, $C_H^o(\pi_1(y_H))$, is higher than expected consumption in period 1.

For low values of $\pi_1(y_H)$, condition (14) holds, the lower bound on the value of high-income consumers is high, and the incumbent's profits in period 2 are too low to attain expected consumption equalization. Thus, the best that can be done is to maximize such profits, and the optimal way to do so is by providing no information.

For higher values of $\pi_1(y_H)$, condition (14) is not satisfied and the disclosure policy is designed to perfectly smooth consumption between period 1 and period 2 after a high income realization in period 1,

$$C_1 = C_2(y_H). \tag{15}$$

This can be achieved by considering two signals only, $M = \{g, b\}$ (good or bad), with a *bad-news* structure: All high-income consumers, together with a fraction of low-income consumers, receive a good signal. Only low-income consumers receive a bad signal. It is optimal to have some low-income consumers with a good signal—even if, as we have shown, this does not affect their consumption—to manipulate the value for the high-income consumers and equalize their consumption between period 1 and period 2. This is attained for an intermediate value of signal informativeness between full information revelation and no information disclosure.

The following proposition characterizes the optimal information disclosure:

Proposition 2. *There exists a cutoff π^* that solves (14) with equality such that:*

1. *If $\pi_1(y_H) < \pi^*$, then $C_1 < C_2(y_H)$ and it is optimal to provide no information;*

2. *If $\pi_1(y_H) \geqslant \pi^*$, then consumption is equalized between periods 1 and 2 after a high income realization, and the optimal disclosure policy has a bad-news structure, i.e., $M = \{g, b\}$ (good or bad), and $\mu(g|y_H) = 1$ and $\mu(b|y_L) \in (0, 1)$, so a bad signal fully reveals period 1 income.*

First, due to the convexity of $C_H^o(s)$, it is optimal to assign the same signal to all high-income consumers. This is because introducing dispersion for the high-income consumers simply increases the cost of providing utility to such consumers. Thus, without loss of generality, we can consider two signals, $M = \{g, b\}$, and have $\mu(g|y_H) = 1$. We are only left to choose the fraction of consumers with low income in period 1 that receive the

Figure 2: Optimal share of high-income consumers conditional on receiving a good signal

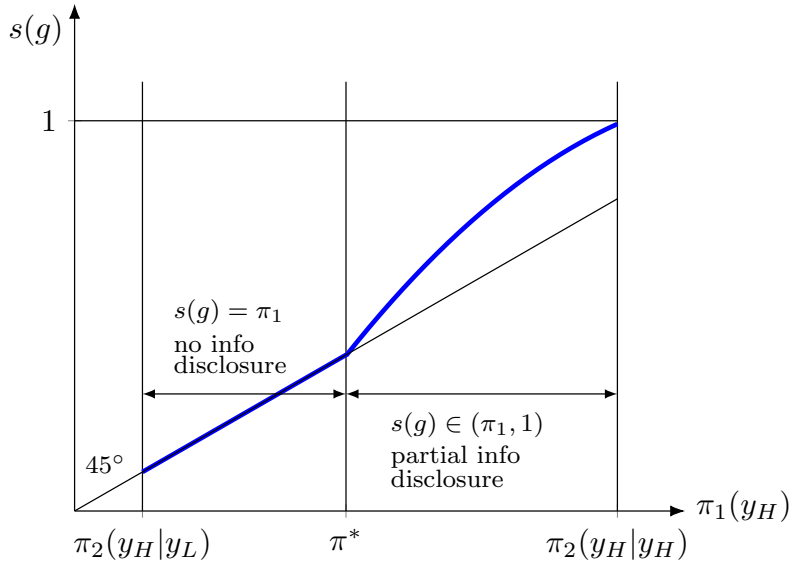

good signal. This is not going to affect their consumption, but it affects the composition of the pool of agents with a good signal and therefore what the incumbent must offer to retain the high-income consumers. The information design problem can then induce any continuation values for the high-income consumer in the range $[V_H^o(\pi_1(y_H)), V_H^o(1)]$.[14]

We can then write the problem as

$$\max_{C_{2H}} u(C_1(C_{2H})) + \pi_1(y_H)\, u(C_{2H}) + \pi_1(y_L)\, u(Y_{2L}) \tag{16}$$

subject to

$$C_1(C_{2H}) = Y_1 + \pi_1(y_H)[Y_{2H} - C_{2H}]$$

$$C_{2H} \in [C_H^o(\pi_1(y_H)), Y_{2H}]$$

and then recover $s(g)$ and $\mu(b|y_L)$ from

$$C_{2H} = C_H^o(s(g)) \text{ and } s(g) = \frac{\pi_1(y_H)}{\pi_1(y_H) + \pi_1(y_L)(1 - \mu(b|y_L))}. \tag{17}$$

The optimal disclosure policy can then be represented by the share of high-income

[14]The size of the interval $[V_H^o(\pi_1(y_H)), V_H^o(1)]$ is hump-shaped in the degree of income persistence. If income is very persistent, e.g. $\pi_2(y_H|y_H) \to 1$, then it is easy to separate the high-income consumers, and $V_H^o(\pi_1(y_H))$ is close to $V_H^o(1)$. The same is true if income is (close to) i.i.d. over time and y_1 is not informative about y_2. Thus, the optimal public disclosure policy is more effective for intermediate values of persistence.

consumers who received a good signal, $s(g)$. In the appendix, we show that for low levels of $\pi_1(y_H)$, condition (14) holds, so that the constraint $C_{2H} \geqslant C_H^o(\pi_1(y_H))$ is binding at the optimum. It is therefore optimal to provide no information, $s(g) = \pi_1(y_H)$, as illustrated in Figure 2. This is because even by maximizing period 2 profits on high-income consumers, by minimizing their values in equilibrium, we still have $C_1 < C_{2H}$.

For higher values of $\pi_1(y_H)$, the value of expected profits under no information disclosure is higher (as is Y_1), so condition (14) is not satisfied and the C_{2H} that solves (16) is interior.[15] It follows that

$$C_1 = C_{2H} = \frac{Y_1 + \pi_1(y_H) Y_{2H}}{1 + \pi_1(y_H)}$$

and $s > \pi_1(y_H)$, as shown in Figure 2. Moreover, since in this region C_{2H} is increasing in $\pi_1(y_H)$ and $V_H^o(s)$ is increasing in s, then s must be increasing in $\pi_1(y_H)$.

Consider now the expected consumption profile at the optimal public disclosure policy. Expected consumption is *front-loaded* and, for high values of $\pi_1(y_H)$, there is perfect insurance between period 1 and period 2 after $y_1 = y_H$:

$$C_1 = C_2(y_H) > C_2(y_L). \tag{18}$$

This is the opposite consumption configuration relative to the standard case with firm commitment (e.g., Harris and Holmstrom (1982) and Thomas and Worrall (1988)). In such a case, consumption is back-loaded and there is perfect insurance between period 1 and period 2 after a bad income realization. The different predictions for the optimal consumption profiles are illustrated in Figure 3. The critical difference is that when the incumbent can commit, it redistributes the profits it can extract from the high-income consumers in period 2 to consumption in period 1 and to low-income consumers in period 2. Consumption smoothing may not be perfect, because the ex-post participation constraint for the high-income consumer may be binding, but consumption is constant otherwise. Without commitment on the firm's side, the incumbent does not provide any transfers to the low-income consumers. (Nor do the outsiders: because of adverse selection, they anticipate that only low-income consumers will switch in equilibrium.)

Most surprisingly, the larger the share of high-income consumers in the economy, $\pi_1(y_H)$, the more likely it is that consumption is front-loaded across periods and that it is optimal to reveal (some) information. This is because, absent information revelation, adverse selection implies that consumption in period 1 increases *more* than consumption

[15] For values of $\pi_1(y_H)$ higher than $\pi_2(y_H|y_H)$ violating our requirement that $Y_1 \leqslant Y_{2H}$, condition (14) fails but it is not possible to find a disclosure policy such that consumption in period 1 and 2 are equated. Thus, in this case it is optimal to provide full information disclosure and consumption is $C_1 = Y_1 > Y_{2H} = C_2(y_H)$. Note that the requirement that $Y_1 \leqslant Y_{2H}$ is never violated in the type interpretation of our economy.

Figure 3: Consumption profile under optimal disclosure policy for high values of $\pi_1(y_H)$

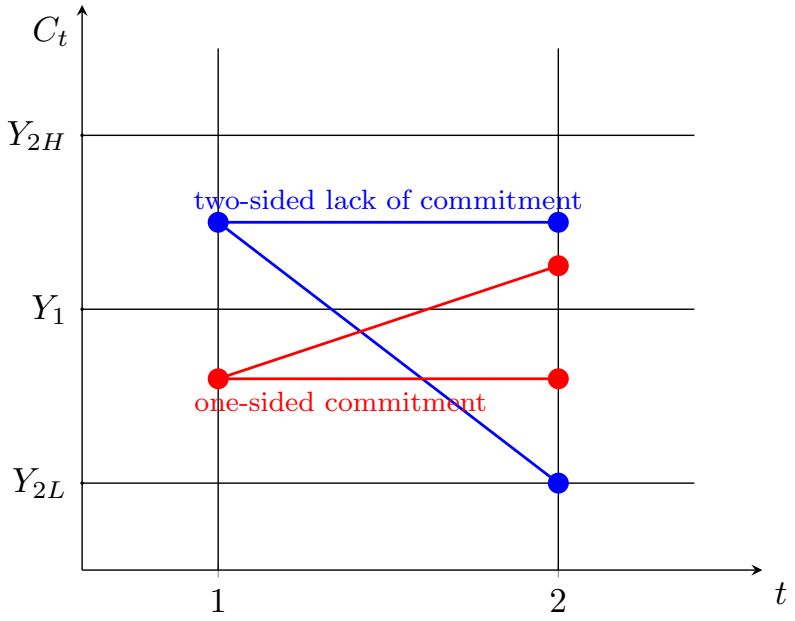

in period 2 for high-income consumers as $\pi_1(y_H)$ grows larger. This result crucially depends on the dynamic nature of the contract, and it delivers the opposite relationship between pool composition and information disclosure compared with a traditional static adverse selection economy, as we show in Section 7 below.

The time series of contractual terms can thus help disentangle whether firms can commit. As we showed in Proposition 1 and Proposition 2, the ability of firms to commit to a contract (because of reputation considerations, for example) is a critical factor for the optimal amount of information to disclose. Front-loading of consumption suggests firms' inability to commit and more information must be provided. Back-loading of consumption suggests that firms can commit and less (no) information should be disclosed to competitors.

6 Discussion

We now discuss some of the assumptions of our model and some extensions.

Regulation and commitment We model the choice of the optimal public disclosure as one made by a regulator. The presence of a regulator is not needed for the implementation of the optimal public disclosure policy. A firm in period 1 endowed with a commitment

technology for reporting information in period 2 will choose the same disclosure policy as the planner to maximize the consumers' welfare subject to the zero-profit condition.

A commitment technology to truthfully reporting according to (M, μ) is necessary. Consider a case in which condition (14) does not hold and it is optimal to provide information. The commitment technology is needed because in period 2 the incumbent would have incentives to provide no information to maximize its profits. Our required notion of commitment is stronger than the one in Lin and Liu (2024), who impose that the distribution of signals ex-post must be conforming with μ^*. This is because, at our equilibrium, the incumbent will have incentives ex-post to assign the good signal to all low-income consumers—at least as long as it is feasible and not detectable—and a bad signal to the high-income consumers. Thus, providing some information is not credible without a commitment technology.

Ownership of information We implicitly assume that firms own the information they observe and are the only ones to have access to a commitment reporting technology that allows them to truthfully communicate, according to (M, μ). Clearly, if consumers had access to such a technology in the first period, they would choose the same disclosure policy as the one in the paper. However, if consumers had access to a technology to credibly report information about y_1 at the end of period 1—but they could not commit to a reporting strategy—then full information disclosure would be the only equilibrium outcome, following the logic in Grossman (1981) and Milgrom (1981). In fact, all high-income consumers would have an incentive to fully disclose y_H.

Allowing firms to trade data does not substitute for a disclosure policy. Suppose the incumbent can sell its data ex post to a single outside firm, committing not to use it afterward so that =ownership of the data maintains positive value. The buyer then earns the same profits on high-income consumers as the incumbent earns under no disclosure. Competition among buyers drives the price of the data up to these profits, and period-1 competition rebates the price to consumers, so the allocation coincides with the one under no information disclosure. When the pool of consumers is sufficiently good, this allocation is dominated by the equilibrium under mandated (partial) disclosure.

Connection with information design literature Our environment departs from the traditional information design literature (Kamenica and Gentzkow (2011), GK) along at least three dimensions. First, it is dynamic, in that information in the second period determines the incumbent's profits that are rebated to the consumer in the first period. Second, given an information structure, payoffs are determined by the equilibrium in a game between

multiple agents. Third, due to both risk aversion and the dynamic nature of the problem, the objective function of the designer/sender typically depends on the whole posterior distribution and not just the posterior mean. Hence, the particular characteristics of our environment do not allow us to cast the information design problem as a variation of the well-known GK framework and its associated graphical representation.

Market structure The design of the optimal disclosure policy hinges on the fact that the share of high types with a given signal, s, affects their outside option in equilibrium, $V_H^o(s)$. This is true in the strategic equilibrium notion that we adopt following Netzer and Scheuer (2014), Hellwig (1987), Wilson (1977), and Miyazaki (1977). Our results remain valid in a variant where outsiders are restricted to a single pooling contract, as in the labor market application of Kahn (2013). By contrast, the conclusions would change under the equilibrium notions of Azevedo and Gottlieb (2017), Bisin and Gottardi (2006), or the competitive search economy of Guerrieri et al. (2010), in which the second-period outside option absent the informed incumbent is always the least-cost-separating allocation and does not depend on s. There the optimal disclosure problem collapses to a binary choice between full information and no information.

Extensions We consider four extensions, all formally presented in the Appendix. In the first, the period-1 outcome is public but the incumbent and the consumer share a private signal about future income (Appendix C). Because the incumbent still earns profits on consumers with a favorable private signal, optimal disclosure is again driven by intertemporal smoothing. What drives our results is the incumbent's informational advantage, not the fact that outsiders do not observe any outcome in period 1. In the second, we allow the incumbent to offer different contracts to consumers with the same history (Appendix D). The incumbent then offers most high types strictly less than the outside option they would face under equal treatment, exploiting the fact that any outsider who undercuts it would attract a worse pool of switchers. This extracts more surplus and front-loads consumption further. The optimal disclosure policy still takes the same bad-news form, and it is weakly more informative. In the third, the consumer exerts unobservable effort in the first period in order to draw from a better income distribution in period 2 (Appendix E). Also here the incumbent discloses weakly more information than in the benchmark, because a wider spread in continuation values induces the consumer to exert effort. In the fourth, the consumer can (privately) save between periods at a gross rate equal to the intertemporal discount rate (Appendix F). Savings allow consumers to raise consumption after a low income realization, so revealing no information, hence maxi-

mally front-loading the transfer from the incumbent, becomes optimal for every value of $\pi_1(y_H)$.

7 Comparison with static contracts

We briefly characterize the optimal disclosure policy in a static benchmark: a one-period economy in which firms are symmetrically uninformed about the consumer's type. This benchmark is the natural reference for the equilibrium period 2 allocation in both applications that follow. In Section 8, in the context of consumers switching to more attractive entrants, it emerges as the limiting case in which the entrants' advantage is strong enough that all consumers leave the incumbent. In Section 9, it describes the community-rating regime in health insurance, where the no-discrimination requirement strips the incumbent of its informational advantage. In both cases, the absence of asymmetric information across firms restores the possibility of transfers to low-income consumers.

For a given public disclosure policy, consumption in the first period is Y_1, as there are no future expected profits. In the second period, the spot contract offered by the outsiders to consumers with a signal with share s is the Miyazaki-Wilson contract that solves (2). The expected value without long-term relationships is then equal to

$$u(Y_1) + \sum_s p(s) W(s),$$

where $p(s)$ is the distribution of high types' shares induced by the disclosure policy and

$$W(s) \equiv sV_H^o(s) + (1-s) V_L^o(s)$$

is the expected value associated with the Miyazaki–Wilson contract. Recall that with an informed incumbent, high-income consumers obtain $V_H^o(s)$ in period 2 while low-income consumers obtain only $u(Y_{2L})$ rather than $V_L^o(s)$.

The value W is illustrated in Figure 4 together with V_H^o and V_L^o. As shown in Lemma 1, for $s < s^*$ the value for the low-income consumers is $u(Y_{2L})$ and for the high-income consumers it is V_H^{LCS}. Both are independent of s, so W is linear in s over the interval $[0, s^*)$. For shares above s^*, the solution to (2) calls for subsidizing the low types to relax the incentive compatibility constraint. When that happens, both types receive the same (undistorted) consumption allocation, which is equal to $sY_{2H} + (1-s) Y_{2L}$.

Figure 4: Expected continuation value with spot contracts, $W(s)$

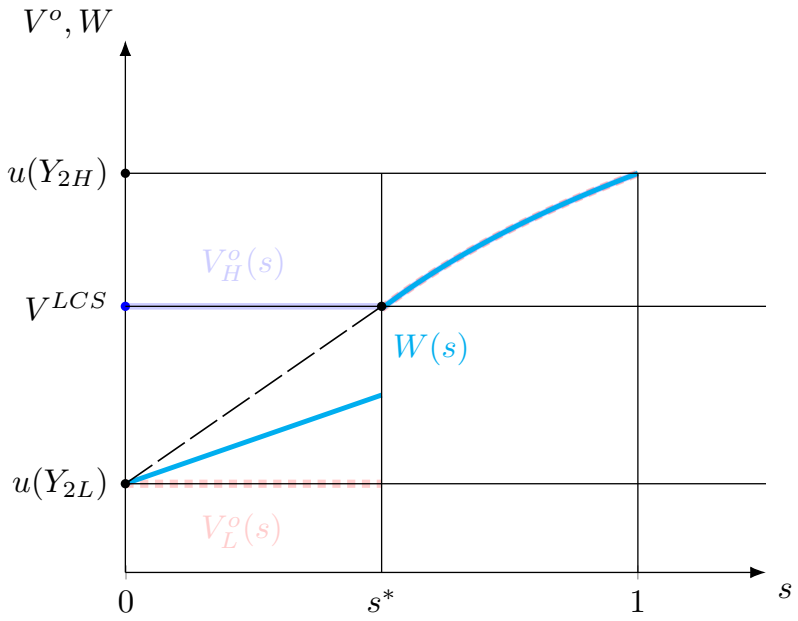

The optimal disclosure policy solves

$$\max_{p(s)\in\Delta([0,1])} \sum_s p(s) W(s) \tag{19}$$

subject to $\sum_s p(s)\, s = \pi_1(y_H)$.

Lemma 6. *The optimal disclosure policy is such that:*

1. *If $\pi_1(y_H) < s^*$, then it is optimal to provide partial information disclosure with a bad-news structure: $M = \{g, b\}$ with $s(g) = s^*$ and $s(b) = 0$.*

2. *If $\pi_1(y_H) \geqslant s^*$, then it is optimal to provide no information.*

The optimal information design problem trades off the cross-subsidization it delivers to the low-income consumers and the distortions associated with the incentive constraint for the high-income consumers. If $\pi_1(y_H) < s^*$, it is optimal to provide partial information disclosure. The share of the high type among the consumers with a good signal is s^*. Mechanically, the partial information disclosure takes advantage of the non-concavity of W over $[0, s^*]$, caused by the upward jump at s^*, and has two signals with shares 0 and s^*.

Intuitively, if $\pi_1(y_H)$ is low, the composition of the consumer pool absent any information disclosure is bad enough that it is optimal for the outsiders to offer a heavily distorted

allocation to the high-income consumers and no subsidies to the low-income consumers. Thus, disclosing some information is beneficial to both the high-income consumers that can attain a less distorted allocation and, in expectations, also to the low-income consumers. In fact, with some disclosure, at least a fraction of the low-income consumers receives a subsidy. This result is similar to that in Goldstein and Leitner (2018).

When $\pi_1(y_H) > s^*$, it is optimal to provide no information because the equilibrium contract features pooling. Under information disclosure, the lower distortions for the high-income consumers are dominated by the reduction in insurance for the low-income consumers.

Lemma 6 implies that, regardless of $\pi_1(y_H)$, at least a fraction of the low-income consumers receive consumption higher than Y_{2L}. Thus, the absence of an informed incumbent has the advantage of delivering higher consumption for the low-income consumers in the second period. There are also costs: no intertemporal smoothing is possible, and there might be distortions in the allocation offered to the high-income consumers in the second period in order to screen out the low-income ones. The next two sections will highlight the trade-off between such costs and benefits in the contexts of lending and health insurance contracts.

8 Application 1: Open Banking

The model offers insights into any market in which an incumbent firm acquires an informational advantage over its competitors. Our first application is the bank lending market, whether corporate loans or personal credit.

This application is particularly relevant for two reasons. First, recent work provides direct evidence that incumbent banks are better informed about a borrower's creditworthiness than the rest of the market. For example, Beyhaghi et al. (2024) use regulatory data on banks' internal assessments of corporate loan default probabilities and losses given default, and find that those assessments predict future stock and bond returns as well as earnings surprises relative to analyst forecasts and to predictions based on publicly available data. Their results are stronger for small firms but hold for public companies as well.

Second, our analysis speaks directly to the recent policy debate about *open banking*. Under the EU's open banking directive—and similar proposals under discussion in the United States—incumbent financial firms must, at the customer's request, share client information with competing financial firms. The consequences of open banking on bank competition and entry have been studied in recent empirical (Babina et al. (2024)) and

theoretical (He et al. (2023)) work. Both papers restrict attention to static contracts, hence neglect intertemporal smoothing and its interaction with cross-subsidization. Our optimal disclosure policy can then inform the debate about how much information banks should be forced to disclose.

To map our model to the lending environment, assume that consumers are borrowers and firms are lenders. The contract between a borrower and lender is (b, R), where b is the loan size and R is the gross interest rate charged. Borrowers repay the loan if they have a high-income realization, y_H, otherwise the bank receives a recovery value y_L. Let

$$b = c_{am} \text{ and } R = \frac{y_H - c_{pm}}{b}.$$

The zero-profit condition then implies that the expected gross return on the loan equals one—statically for the contracts offered by outsiders in period 2, and in present value for the incumbent.

Under two-sided lack of commitment, the model predicts that the interest rate offered in the contract starts low and eventually increases: the bank subsidizes the initial interest rate to compensate for the profits it makes in the second period on the high-income borrowers. This is consistent with Ioannidou and Ongena (2010), who show that new loans carry lower interest rates and that banks subsequently raise the cost of lending to firms with whom they already have a relationship.

We finally turn to the implications for open banking. Our analysis contains two main messages. First, unrestricted open banking, in the sense of full information sharing, is never optimal. Importantly, this conclusion holds even if open banking requires borrowers to opt in, as in the majority of proposals being discussed. Thus, assigning property rights over information to consumers is not optimal. Second, when banks lack commitment, a partial form of open banking can be welfare-improving, in particular in markets in which the typical firm repays and only a few default (high π_1), as shown in Proposition 2. Absent any information disclosure, high values of π_1 trigger excessive front-loading of loan quantities, and interest rates that rise too much within the relationship.

Extension: Switching Lenders

The bank-lending application above gave borrowers no allocative motive to switch lenders: outsiders attracted consumers only via the terms of the contract offered. In practice, however, fintech entrants—algorithmic underwriters, digital-only banks, embedded-finance platforms—offer borrowers additional advantages such as faster origination, lower transaction costs, or broader access. We capture this by extending the baseline so that outsiders

deliver an additional utility benefit $\Delta \geqslant 0$ to consumers who switch. Creditworthy borrowers may then switch in equilibrium, which ameliorates the adverse-selection problem at the outsider and supports cross-subsidization in the second period.

Concretely, a consumer who switches to a new firm enjoys utility $u(C)+\Delta$, where C is the expected per-period consumption prescribed by the outsider's contract. This additive form preserves the Miyazaki-Wilson contract described in (2), so $V_i^o(s;\Delta) = V_i^o(s) + \Delta$ for $i = H, L$.

Consider the period-2 equilibrium for a given disclosure policy. When $\Delta > 0$, all low-income consumers leave the incumbent—new firms offer them at least $u(Y_{2L}) + \Delta$, the incumbent at most $u(Y_{2L})$—breaking the indeterminacy of the baseline with $\Delta = 0$. High-income consumers also switch when their signal's pool share is high enough: the incumbent retains them only if $V_H^o(s) + \Delta \leqslant u(Y_{2H})$, that is, only for s below the cutoff $\hat{s}(\Delta)$ implicitly defined by $V_H^o(\hat{s}(\Delta)) + \Delta = u(Y_{2H})$. Throughout we restrict attention to $\Delta \in [0, \bar{\Delta}]$, where $\bar{\Delta}$ is the benefit level at which high types strictly prefer to switch even at the least-cost separating contract, $V_H^{LCS} + \bar{\Delta} = u(Y_{2H})$. For larger Δ the allocation coincides with that at $\bar{\Delta}$.

Switching at signals $s \geqslant \hat{s}(\Delta)$ has two consequences. The incumbent makes no period-2 profit on these consumers, so period-1 consumption is correspondingly lower. Conversely, the resulting improvement in the outsiders' pool relaxes adverse selection and lets low-income consumers earn a payoff above $u(Y_{2L}) + \Delta$ whenever $s \geqslant s^*$ (so the Miyazaki-Wilson contract calls for cross-subsidization). It is easy to see that $s^* \leqslant \hat{s}(\Delta)$ whenever $\Delta \in [0, \bar{\Delta}]$, so high-type switching always coincides with cross-subsidization of the low types.[16]

The optimal disclosure policy trades off the intertemporal smoothing of the baseline ($\Delta = 0$) against the period-2 cross-subsidization and allocative gains that switching makes available. Let $W(s;\Delta)$ denote the expected period-2 continuation value at posterior s, where low-income types' value drops from $V_L^o(s)$ to the autarky payoff $u(Y_{2L})$ once s falls below $\hat{s}(\Delta)$. The optimum admits a simple three-signal characterization:

Lemma 7. *The optimal disclosure policy can be implemented with $M = \{b, i, g\}$ (bad, intermediate, good), where b is a bad signal and $s(b) = 0$, $s(i) \equiv s_{in} \in (0, \hat{s}(\Delta))$, so high-income consumers with signal i stay with the incumbent, and $s(g) \equiv s_{out} \in [\hat{s}(\Delta), 1]$, so high-income consumers with signal g switch to an outsider.*

Only one switching signal is needed—the incumbent makes no profit on switchers and W is concave above $\hat{s}(\Delta)$—and the logic of Proposition 2 delivers the analogous result for

[16] To see this, combining the definition of $\hat{s}(\Delta)$ and $\bar{\Delta}$, we have $V_H^o(\hat{s}(\Delta)) + \Delta = u(Y_{2H}) = V_H^{LCS} + \bar{\Delta}$. Thus, it must be that $V_H^o(\hat{s}(\Delta)) - V_H^{LCS} = \bar{\Delta} - \Delta > 0$. Since $V_H^o(s)$ is increasing, it follows that $\hat{s}(\Delta) > s^*$.

the stayer signal. The planner then chooses the signal probabilities (p_0, p_{in}, p_{out}) and pool compositions (s_{in}, s_{out}) to maximize period-1 utility plus expected continuation value subject to Bayes plausibility; the explicit program is in Appendix B.

Proposition 3. *Full information is never optimal. Moreover, there exist cutoffs* $0 < \Delta_L < \Delta_H < \bar{\Delta}$ *such that for all* $\Delta \leqslant \Delta_L$*, we have* $p_{out} = 0$*; for all* $\Delta \geqslant \Delta_H$*, we have* $p_{in} = 0$*.*

Proposition 3 identifies two regimes. For small Δ, the incumbent retains all high-income consumers ($p_{out} = 0$); the disclosure problem reduces to the baseline case where the optimal disclosure is driven by intertemporal smoothing considerations, and the optimal signal becomes less informative as Δ grows. This is because the gains for letting the high-income consumers get the utility benefit Δ and allowing a negligible measure of low-income consumers to obtain some cross-subsidization are limited when Δ is small. To maximize the gains from consumption smoothing, it is optimal to keep the continuation value for the high type constant as Δ changes, $V_H^o(s_{in}(\Delta)) + \Delta = V_H^o(s_{in}(0))$. Thus, $s_{in}(\Delta)$ must be decreasing in this range, and the signal for the high type must be less informative, to reduce the value of the outside option and offset the increase in value given by the higher Δ.

For large Δ, the higher utility offered by the outsiders and the subsidy offered in equilibrium to the low type in period 2 outweigh the possible gains from intertemporal consumption smoothing. At the optimum, all consumers switch ($p_{in} = 0$), the incumbent makes no period-2 profits, and the problem collapses to a static one that allows cross-subsidization through pooling.

Figure 5 traces how optimal disclosure responds to the entry of more attractive outsiders—a rise in Δ from Δ_L, close to 0, to Δ_H, close to the upper bound $\bar{\Delta}$—for a fixed pool composition $\pi_1(y_H)$. The key result is that the sign of the response depends on the quality of the pool. When the pool is good (high $\pi_1(y_H)$), the incumbent discloses some information at $\Delta = 0$ to smooth consumption intertemporally, as in our benchmark model. As Δ rises, disclosure should *decrease*, since pooling the many high types who now switch with the low types is what sustains cross-subsidization. When the pool is poor (low $\pi_1(y_H)$), the pattern reverses: no information is disclosed at $\Delta = 0$, but as Δ rises disclosure should *increase*, as only by revealing information can the outsiders be induced to offer a contract that cross-subsidizes the low types. At intermediate Δ the two motives coexist: the optimum uses two informative signals, an intermediate one that keeps some high types with the incumbent for intertemporal smoothing and a top one that sends the rest to outsiders, where they cross-subsidize the low types.

Full disclosure is optimal in neither regime: it destroys both the period-2 rents that finance intertemporal smoothing and the pooling that sustains cross-subsidization. This

Figure 5: How information content varies by Δ for different consumer pool compositions

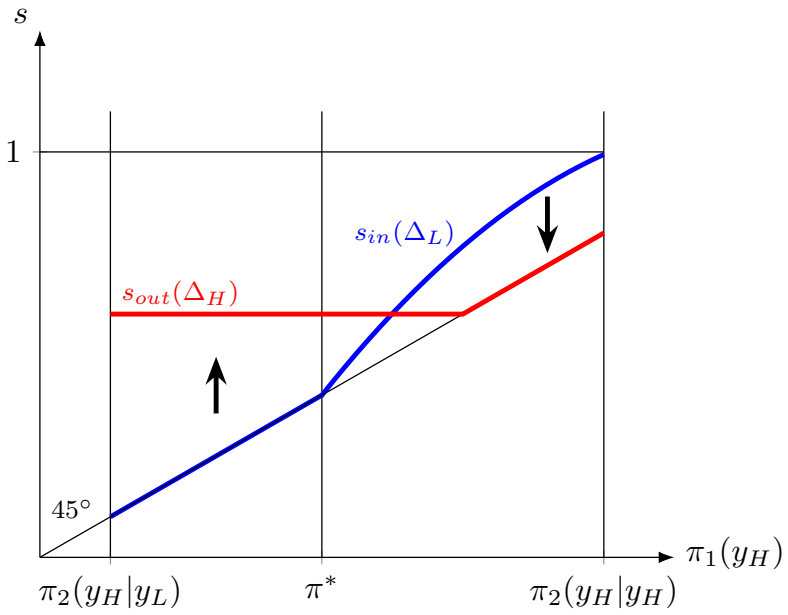

result departs from Mukherjee (2008), who finds full information optimal in environments with firm commitment and transferable rights (exit fees or severance payments) that allow the incumbent to extract switchers' surplus and rebate it ex ante.[17] This is not the case when exit fees, severance payments, or both are not available, or firms' commitment is not possible.[18] Without commitment, changes in information disclosure affect not only the amount of resources available to the incumbent but also the way they are distributed. The additional resources extracted from high-income consumers cannot be redistributed to the state in which they are most valued. For sufficiently high s, reducing such misallocation of resources always dominates with respect to increasing their total quantity, making full information suboptimal.

The switching extension above yields additional implications for open banking. First, the pool of switchers is typically worse than the pool of stayers due to adverse selection.[19] This result echoes the finding in the personal loan market by Di Maggio and Yao (2021)

[17]Full information disclosure attains the same allocation as in Cochrane (1995), where health status is public information and severance payments implement full insurance.

[18]One can also show that with only commitment or only exit fees, full information is never optimal. In the first case, full information lowers the resources available to be distributed to low-income consumers. In the second case, full information eliminates cross-subsidization by the only firms that provide it, i.e., outsiders.

[19]For example, in the equilibrium characterized above, all low-income borrowers switch if $\Delta > 0$ while only a subset of high-income borrowers do.

who document higher default rates among borrowers that switch to a fintech firm, but decreasingly so with time spent by the fintech firm in the market. Moreover, hard information like credit score and loan size explains 80% of the variation in interest rates for fintech loans but only 59% for bank loans.

Second, when new technologies emerge that deliver additional value to customers, information sharing should be reduced when the pool of borrowers is sufficiently good. When the introduction of new technologies induces creditworthy borrowers to move away from the incumbent, the adverse-selection problem is less severe and it is possible to cross-subsidize less creditworthy borrowers by providing less information to the new firms. This is valuable in our economy because borrowers are risk-averse and prefer an allocation that insures them against low income realizations. Alternatively, the same result would arise in an economy where the government has redistributive motives.

9 Application 2: Health Insurance

Our model naturally applies to health insurance markets, where an incumbent insurer accumulates private information about a policyholder's health over time while competitors are limited to coarser signals such as community-rating classifications. This application is policy-relevant for two reasons. First, a long literature emphasizes the *reclassification risk*, that is, the risk of raising insurance premiums after a negative health shock, that arises as insurers learn about their policyholders over time (Handel et al. (2015)). Second, our analysis speaks directly to the policy debate about community rating and risk classification in the U.S. individual market, which has gained urgency after the 2017 repeal of the Affordable Care Act's individual mandate to purchase coverage.

To map the model to a health insurance application, let firms be insurers, and consumers be policyholders. The labels am and pm now refer to health states rather than subperiods: am is the bad-health state, while pm is the good-health state, which occurs when $y = y_H$. Given our transition matrix, a policyholder who is healthy in period 1 is more likely to remain healthy in period 2. An insurance contract specifies a consumption pair $(c_{t,am}, c_{t,pm})$, where $c_{t,am}$ is consumption in the bad-health state and $c_{t,pm}$ is consumption in the good-health state. The contract can equivalently be expressed in terms of coverage X_t and a premium P_t. Consumption in the two health states is then $c_{t,am} = X_t$ and $c_{t,pm} = y_H - P_t$. A higher coverage therefore provides greater protection in the bad-health state, whereas a lower premium raises consumption in the good-health state. Utility within the period is aggregated as in the general model. Since $\beta < 1$, the lower weight on good-health-state consumption captures the idea that resources are more

valuable when the policyholder is in bad health. The welfare-maximizing allocation sets $P_t = y_H$ and $X_t = (Y_1 + Y_2)/2$.

Under two-sided lack of commitment, the model predicts that period-1 premiums are subsidized by the insurer's expected future profits, and premiums rise in period 2 as the insurer accumulates information about the policyholder's health—sharply so for those revealed to be unhealthy, who face less favorable terms going forward. This is the reclassification dynamic that the literature on long-term health insurance has emphasized (Handel et al. (2015)).

The incumbent's competitors must screen policyholders whose history they cannot observe: the separating contract intended for those with a healthy history features a lower premium and lower coverage—a distortion that healthy policyholders tolerate precisely because they are unlikely to need the coverage—while those with a costly history obtain full coverage at a higher premium. This is the classic pattern of insurance markets under adverse selection, and in our dynamic setting its severity is governed by the disclosure policy through the composition s of each rating class.

We turn to the implications for the design of risk classification. Our analysis contains two main messages. First, full disclosure of a policyholder's health history to outsiders is never optimal—even in this experience-rating regime where reclassification is permitted: it eliminates the period-2 informational rents that competition rebates as period-1 premium discounts, destroying intertemporal smoothing. Second, when the pool of policyholders is sufficiently healthy (high $\pi_1(y_H)$), a partial form of disclosure—a coarse risk-classification system that pools some sick policyholders with healthy ones—is welfare-improving, as shown in Proposition 2.

Extension: Experience vs. Community Rating

So far we have studied health insurance under the baseline experience-rating regime, in which the incumbent uses its period-1 health information to price the period-2 contract. We now consider the alternative regime—community rating—in which the insurer is prohibited from conditioning the period-2 contract on the period-1 health realization y_1.

The Affordable Care Act's community rating rules prohibit insurers from conditioning premiums on a policyholder's health history, regardless of what the incumbent observes. Community rating restricts the period-2 premium P_2 and coverage X_2 to be independent of the period-1 health realization y_1. The incumbent is then forced to offer menus satisfying incentive compatibility despite its informational advantage, implementing the spot-contract equilibrium of Section 7.

This restriction has been the subject of sustained policy debate in recent years. The 2017 Tax Cuts and Jobs Act effectively repealed the individual mandate by zeroing out its penalty starting in 2019, and successive administrations have considered further reforms—including the expansion of short-term, limited-duration plans that compete with ACA-compliant marketplaces. Both academic and policy work have raised concerns that, absent a mandate, the consumer pool in the individual market may deteriorate as healthier policyholders exit, weakening the pre-conditions under which community rating sustains insurance.[20] Hackmann et al. (2015) document the role of an individual mandate in reducing adverse selection in the Massachusetts pre-ACA reform. Saltzman (2019) estimates that the mandate repeal substantially raises premiums and reduces enrollment. These developments motivate revisiting when community rating is the right design and what shape experience rating should take when optimal.

Community rating corresponds to the static benchmark of Section 7. The baseline experience-rating regime corresponds to the economy with long-term relationships analyzed in Section 5. The next proposition shows that the optimal regime depends on the quality of the pool: for sufficiently low $\pi_1(y_H)$ (an unhealthy pool) experience rating dominates, while for $\pi_1(y_H)$ sufficiently high (a healthy pool) community rating dominates.[21]

Formally, note that our stated assumptions require that $0 < \pi_2(y_H|y_L) \leqslant \pi_1(y_H) \leqslant \pi_2(y_H|y_H) < 1$ to ensure that $Y_{2L} \leqslant Y_1 \leqslant Y_{2H}$ and that at least one inequality is strict so that there is still residual uncertainty in period 2.

Proposition 4. *If $\pi_1(y_H)$ is sufficiently close to the lower bound $\pi_2(y_H|y_L)$, ex-ante welfare is higher under experience rating (the economy with long-term relationships) than under community rating (the economy with spot contracts only). For $\pi_1(y_H)$ sufficiently close to the upper bound $\pi_2(y_H|y_H)$, the ranking is reversed.*

Intuitively, when there are sufficiently few high-income consumers, in the economy with spot contracts only, the second-period allocation entails little cross-subsidization across consumers. However, due to its informational advantage, the incumbent is able to avoid distorting the consumption of the high-income consumers and to distribute the ex-post profits $Y_{2H} - C\left(V_H^{LCS}\right)$ as first-period consumption. This intertemporal smoothing is very valuable when the pool in the first period is not good because consumption

[20] See Geruso and Layton (2017) for a survey of selection in health insurance markets and its policy remedies.

[21] The comparison is done under the optimal information disclosure for each case. The ranking of regimes would be preserved if the comparison were made without optimizing information disclosure. This is because, for a given pool quality, no information revelation is optimal for the dominant regime, while it is strictly suboptimal for the dominated one.

Figure 6: Value of long-term relationships between firms and consumers

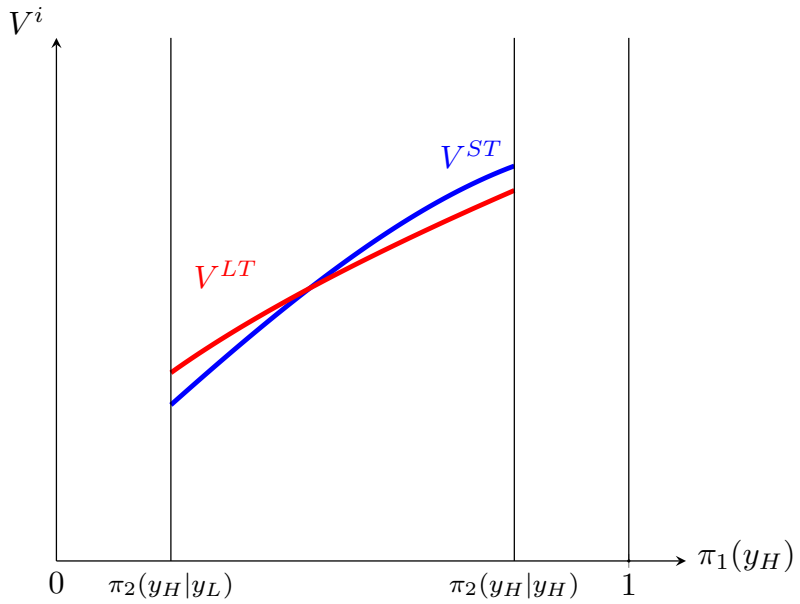

in the first period is lower than consumption in the second period conditional on a good shock as shown in Proposition 2 part 1. It follows that when adverse selection is severe, long-term relationships are beneficial.

Conversely, when the pool of consumers in the economy has sufficiently many high-income consumers, there are large gains from cross-subsidization in the second period, absent the adverse selection induced by the presence of an informed firm. It follows that for a sufficiently large share of high-income consumers, long-term relationships are harmful.

Characterizing the benefit of long-term relationships for interior values of π is more challenging. Yet, in all our numerical examples, we find that the intuition we provided for the ranking of values at the extremes of the support of π holds in the interior as well. Hence, there is a cutoff π above which $V^{\mathrm{ST}}(\pi) > V^{\mathrm{LT}}(\pi)$, as illustrated in Figure 6.

The central trade-off in the existing literature is between adverse selection distortions and reclassification risk, studied by Handel et al. (2015) in a static framework with symmetrically uninformed firms. Community rating protects sick policyholders against premium spikes when their health deteriorates but invites adverse selection as healthy policyholders exit the pool. Experience rating eliminates distortions caused by adverse selection but exposes sick policyholders to higher insurance prices. Their main result is that community rating yields higher ex-ante consumer welfare—expected utility evaluated

before health status is realized—when the pool is sufficiently healthy and the gains from cross-subsidization of sick policyholders are large, while experience rating is preferable when $\pi_1(y_H)$ is low.[22]

Our analysis agrees with Handel et al. (2015) when $\pi_1(y_H)$ is large: Proposition 4 shows that spot contracts with no information being disclosed, hence community rating, dominate long-term relationships when the pool is sufficiently healthy. For low values of $\pi_1(y_H)$, the static benchmark of Section 7 also favors the use of health information, calling for near-full disclosure in line with their prescription. Adding the dynamic dimension, absent from their framework, changes the policy recommendation. Proposition 4 shows that the optimal arrangement is the long-term contract with no information disclosed to the outsiders. The *informed* incumbent, who adopts an experience rating, earns period-2 profits on healthy policyholders when facing *uninformed* outsiders, who are necessarily limited to community rating. Ex-post profits are rebated as lower opening premiums due to ex-ante competition. Sharing information with outsiders would drive those profits to zero and destroy intertemporal consumption smoothing without delivering any gains. To put it concisely, for low values of $\pi_1(y_H)$, consumers want one and only one firm to know their type.

10 Conclusion

We studied optimal public disclosure of information in a dynamic insurance economy in which the incumbent firm has an informational advantage over its competitors. If the incumbent can commit, it is optimal to disclose no information: more information allows outsiders to bid for high-income consumers more aggressively, limiting the transfers available to low-income consumers. If the incumbent cannot commit, second-period competition disciplines its ex-post behavior: no cross-subsidization survives, but partial disclosure governs intertemporal smoothing between consumption in the first period and in the second period after a high income realization. Full information is never optimal, as it erases the ex-post incumbent profits that are rebated in the first period.

We used the model to study two current policy debates—open banking in lending markets and community rating in health insurance—and how optimal disclosure in each case depends on the composition of the consumer pool. In the first application, we asked how credit information sharing should respond to fintech entrants that offer borrowers additional benefits from switching. When most borrowers are creditworthy, some dis-

[22] See also Finkelstein et al. (2009) for evidence on the redistributive implications of policies that restrict the dependency of contracts on observable consumer characteristics.

closure is optimal even absent entry, to implement intertemporal consumption smoothing. As entrants improve, disclosure should decrease, as pooling the many creditworthy switchers with risky borrowers sustains cross-subsidization among switchers. When the pool is poor, the pattern reverses: no disclosure is optimal absent entry, and disclosure should increase as entrants improve, since only a finer credit rating weakens adverse selection and induces entrants to cross-subsidize some risky borrowers. In the second application, the question is whether policyholders prefer that one insurer—or that no insurer—be able to act on their health history. When the pool is healthy, community rating dominates, as suppressing the incumbent's informational advantage restores cross-subsidization of sick policyholders. When the pool is unhealthy, experience rating dominates, but information should not be shared with the market: an informed incumbent facing uninformed entrants avoids the screening distortions of adverse selection and rebates its rents as first-period premium discounts.

We see two main avenues for further research. From a theoretical perspective, our model abstracts from endogenous firm entry and innovation. An information monopoly over a set of locked-in consumers may disincentivize entry and the introduction of technologies that deliver higher utility, as in the case analyzed in Section 8. Modeling entry and studying how it interacts with the provision of insurance is a fruitful area for future research.[23] From an empirical perspective, measuring the model fundamentals in the markets we studied in Sections 8 and 9 would help quantify the forces behind information policy reforms such as open banking and community rating. This is especially valuable because optimal disclosure depends on consumers' incentives to switch and on the composition of the pool, both likely to vary systematically across market segments.

References

ALI, S. N., G. LEWIS, AND S. VASSERMAN (2023): "Voluntary Disclosure and Personalized Pricing," *The Review of Economic Studies*, 90, 538–571. 5

AZEVEDO, E. M. AND D. GOTTLIEB (2017): "Perfect competition in markets with adverse selection," *Econometrica*, 85, 67–105. 22

BABINA, T., S. A. BAHAJ, G. BUCHAK, F. DE MARCO, A. K. FOULIS, W. GORNALL,

[23]This aspect is analyzed by Jin and Vasserman (2021) in the context of the adoption of a monitoring technology in the car insurance market. They show in a policy counterfactual that forcing firms to share data will reduce the incentives to elicit such data and lead to lower welfare in equilibrium. Although via different mechanisms, both their paper and ours show that erasing ex-post profits by fully revealing consumer information is never optimal.

F. MAZZOLA, AND T. YU (2024): "Customer data access and fintech entry: Early evidence from open banking," Tech. rep., National Bureau of Economic Research. 5, 25

BALKE, N. AND T. LAMADON (2022): "Productivity shocks, long-term contracts, and earnings dynamics," *American Economic Review*, 112, 2139–2177. 4, 14

BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): "The limits of price discrimination," *American Economic Review*, 105, 921–957. 4

BEYHAGHI, M., C. HOWES, AND G. WEITZNER (2024): "The Information Advantage of Banks: Evidence From Their Private Credit Assessments," *Available at SSRN 4265161*. 25

BISIN, A. AND P. GOTTARDI (2006): "Efficient competitive equilibria with adverse selection," *Journal of political Economy*, 114, 485–516. 22

BURDETT, K. AND M. COLES (2003): "Equilibrium wage-tenure contracts," *Econometrica*, 71, 1377–1404. 4

COCHRANE, J. H. (1995): "Time-consistent health insurance," *Journal of Political Economy*, 103, 445–473. 29

COHEN, A. (2012): "Asymmetric learning in repeated contracting: An empirical study," *Review of Economics and Statistics*, 94, 419–432. 1

DE GARIDEL-THORON, T. (2005): "Welfare-improving asymmetric information in dynamic insurance markets," *Journal of Political Economy*, 113, 121–150. 4

DI MAGGIO, M. AND V. YAO (2021): "Fintech borrowers: Lax screening or cream-skimming?" *The Review of Financial Studies*, 34, 4565–4618. 29

DOVIS, A. AND P. MARTELLINI (2024): "Long-Term Contracts, Commitment, and Optimal Information Disclosure," Tech. rep., National Bureau of Economic Research. 8, 63

DWORCZAK, P., S. D. KOMINERS, AND M. AKBARPOUR (2021): "Redistribution through markets," *Econometrica*, 89, 1665–1698. 8

EECKHOUT, J. AND L. VELDKAMP (2022): "Data and market power," Tech. rep., National Bureau of Economic Research. 5

ELLIOTT, M., A. GALEOTTI, A. KOH, AND W. LI (2024): "Market segmentation through information," *Available at SSRN 3432315*. 4

FAINMESSER, I. P., A. GALEOTTI, AND R. MOMOT (2023): "Consumer profiling via information design," *Available at SSRN*. 4

FARBOODI, M., P. KONDOR, AND P. KURLAT (2025): "Equilibrium Spillover of Big Data," Working paper. 5

FARINHA LUZ, V., P. GOTTARDI, AND H. MOREIRA (2023): "Risk classification in insurance markets with risk and preference heterogeneity," *Review of Economic Studies*, 90, 3022–3082. 4

FINKELSTEIN, A., J. POTERBA, AND C. ROTHSCHILD (2009): "Redistribution by insurance market regulation: Analyzing a ban on gender-based retirement annuities," *Journal of financial economics*, 91, 38–58. 34

GARCIA, D. AND M. TSUR (2021): "Information design in competitive insurance markets," *Journal of Economic Theory*, 191, 105160. 4

GERUSO, M. AND T. J. LAYTON (2017): "Selection in health insurance markets and its policy remedies," *Journal of Economic Perspectives*, 31, 23–50. 32

GOLDSTEIN, I. AND Y. LEITNER (2018): "Stress tests and information disclosure," *Journal of Economic Theory*, 177, 34–69. 25

GREEN, E. J. (1987): "Lending and the smoothing of uninsurable income," *Contractual arrangements for intertemporal trade*, 1, 3–25. 4

GROSSMAN, S. J. (1981): "The informational role of warranties and private disclosure about product quality," *The Journal of law and Economics*, 24, 461–483. 21

GUERRIERI, V. (2008): "Inefficient unemployment dynamics under asymmetric information," *Journal of Political Economy*, 116, 667–708. 5

GUERRIERI, V., R. SHIMER, AND R. WRIGHT (2010): "Adverse selection in competitive search equilibrium," *Econometrica*, 78, 1823–1862. 5, 22

HACKMANN, M. B., J. T. KOLSTAD, AND A. E. KOWALSKI (2015): "Adverse selection and an individual mandate: When theory meets practice," *American Economic Review*, 105, 1030–1066. 32

HANDEL, B., I. HENDEL, AND M. D. WHINSTON (2015): "Equilibria in health exchanges: Adverse selection versus reclassification risk," *Econometrica*, 83, 1261–1313. 6, 30, 31, 33, 34

HARRIS, M. AND B. HOLMSTROM (1982): "A theory of wage dynamics," *The Review of Economic Studies*, 49, 315–333. 2, 3, 4, 14, 19

HE, Z., J. HUANG, AND J. ZHOU (2023): "Open banking: Credit market competition when borrowers own the data," *Journal of financial economics*, 147, 449–474. 5, 26

HELLWIG, M. (1987): "Some recent developments in the theory of competition in markets with adverse selection," *European Economic Review*, 31, 319–325. 22

HENDEL, I. AND A. LIZZERI (2003): "The role of commitment in dynamic contracts: Evidence from life insurance," *The Quarterly journal of economics*, 118, 299–328. 4, 14

HIRSHLEIFER, J. (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61, 561–574. 2

HOPENHAYN, H. AND I. WERNING (2008): "Equilibrium default," *Manuscript, MIT*, 4, 21. 4

IMMORLICA, N., A. M. SZTUTMAN, AND R. M. TOWNSEND (2022): "Optimal Credit Scores Under Adverse Selection," in *Proceedings of the 23rd ACM Conference on Economics and Computation*, 737–738. 4

IOANNIDOU, V. AND S. ONGENA (2010): ""Time for a change": loan conditions and bank behavior when firms switch banks," *The Journal of Finance*, 65, 1847–1877. 1, 26

JIN, Y. AND S. VASSERMAN (2021): "Buying data from consumers: The impact of monitoring programs in us auto insurance," Tech. rep., National Bureau of Economic Research. 35

JONES, C. I. AND C. TONETTI (2020): "Nonrivalry and the Economics of Data," *American Economic Review*, 110, 2819–2858. 5

KAHN, L. B. (2013): "Asymmetric information between employers," *American Economic Journal: Applied Economics*, 5, 165–205. 1, 22

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian persuasion," *American Economic Review*, 101, 2590–2615. 4, 21, 49

KANIEL, R. AND D. ORLOV (2025): "Intermediated asymmetric information, compensation, and career prospects," *American Economic Review*, 115, 3638–3674. 4

KOVRIJNYKH, N., I. LIVSHITS, AND A. ZETLIN-JONES (2024): "Building credit histories," Tech. rep., Working paper. 5

KURLAT, P. AND F. SCHEUER (2021): "Signalling to Experts," *Review of Economic Studies*, 88, 800–850. 5

LESTER, B., A. SHOURIDEH, V. VENKATESWARAN, AND A. ZETLIN-JONES (2019): "Screening and adverse selection in frictional markets," *Journal of Political Economy*, 127, 338–377. 5

LIN, X. AND C. LIU (2024): "Credible persuasion," *Journal of Political Economy*, 132, 2228–2273. 21

MILGROM, P. R. (1981): "Good news and bad news: Representation theorems and applications," *The Bell Journal of Economics*, 380–391. 21

MIYAZAKI, H. (1977): "The rat race and internal labor markets," *The Bell Journal of Economics*, 394–418. 11, 22

MUKHERJEE, A. (2008): "Career concerns, matching, and optimal disclosure policy," *International Economic Review*, 49, 1211–1250. 4, 5, 29

NETZER, N. AND F. SCHEUER (2014): "A game theoretic foundation of competitive equilibria with adverse selection," *International Economic Review*, 55, 399–422. 2, 9, 10, 11, 22, 46

ROTHSCHILD, M. AND J. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 90, 629–649. 9

SALTZMAN, E. (2019): "Demand for health insurance: Evidence from the California and Washington ACA exchanges," *Journal of Health Economics*, 63, 197–222. 32

SHARPE, S. A. (1990): "Asymmetric information, bank lending, and implicit contracts: A stylized model of customer relationships," *The journal of finance*, 45, 1069–1087. 4

THOMAS, J. AND T. WORRALL (1988): "Self-enforcing wage contracts," *The Review of Economic Studies*, 55, 541–554. 2, 4, 14, 19

WILSON, C. (1977): "A model of insurance markets with incomplete information," *Journal of Economic theory*, 16, 167–207. 11, 22

Appendix

A Model

A.1 Type interpretation

Here we show that the pure exchange economy can be interpreted as one with types. Suppose that a consumer can be one of two types, θ_H and θ_L. The type affects the probability distribution of income. In particular, income in the second subperiod of period 1 and 2 can take on two values, $y_t \in \{y_L, y_H\}$ with $y_H > y_L$. θ_H consumers are more likely to have high income than θ_L consumers, $\pi_L(\theta) = \Pr(y_L|\theta)$ and $\pi_H(\theta) = \Pr(y_H|\theta)$ with $\pi_H(\theta_H) > \pi_H(\theta_L)$. It immediately follows that $\mathbb{E}_H y \equiv \mathbb{E}(y|\theta_H) > \mathbb{E}_L y \equiv \mathbb{E}(y|\theta_L)$.

Both the consumer and the firms do not know the consumer's type at the beginning of the period and learn about it through the observations of the history of income realization. Let $\rho(\theta)$ the common prior of being θ type at the beginning of period 1. After observing a high realization of income in period 1, the prior increases while it decreases after the realization of a low income realization:

$$\rho(\theta_H|y_H) = \frac{\rho(\theta_H)\pi_H(\theta_H)}{\sum_\theta \rho(\theta)\pi_H(\theta)} > \rho(\theta_H), \quad \rho(\theta_H|y_L) = \frac{\rho(\theta_H)\pi_L(\theta_H)}{\sum_\theta \rho(\theta)\pi_L(\theta)} < \rho(\theta_H). \tag{20}$$

Then, the probability that a consumer with history y_1 draws $y_2 = y_H$ is

$$\pi_H(y_1) = \rho(\theta_H|y_1)\pi_H(\theta_H) + \rho(\theta_L|y_1)\pi_H(\theta_L)$$

and the expected income is $\mathbb{E}(y_2|y_1) = \rho(\theta_H|y_1)\mathbb{E}_H y + \rho(\theta_L|y_1)\mathbb{E}_L y$. Thus, expected income in period 2 is higher after a high income realization in period 1 than after a low income realization:

$$\mathbb{E}(y_2|y_H) > \mathbb{E}(y_2|y_L).$$

This economy is equivalent to the one considered in the main text with

$$\begin{aligned}
\pi_1(y_H) &= \rho(\theta_H)\pi_H(\theta_H) + (1-\rho(\theta_H))\pi_H(\theta_L) \\
\pi_2(y_2 = y_H|y_H) &= \rho(\theta_H|y_H)\pi_H(\theta_H) + (1-\rho(\theta_H|y_H))\pi_H(\theta_L) \\
\pi_2(y_2 = y_H|y_L) &= \rho(\theta_H|y_L)\pi_H(\theta_H) + (1-\rho(\theta_H|y_L))\pi_H(\theta_L)
\end{aligned}$$

and

$$\begin{aligned} Y_1 &= \rho(\theta_H)\,\mathbb{E}_H y + (1-\rho(\theta_H))\,\mathbb{E}_L y \\ Y_{2H} &= \mathbb{E}(y_2|y_H) \\ Y_{2L} &= \mathbb{E}(y_2|y_L). \end{aligned}$$

As we describe in the text when discussing Proposition 2, providing full information is never optimal. This is because in this type formulation it must be that

$$Y_1 = \pi_1(y_H)\,Y_{2H} + (1-\pi_1(y_H))\,Y_{2L} = Y_2$$

because $\rho = \pi_1(y_H)\,\rho(\theta_H|y_H) + (1-\pi_1(y_H))\,\rho(\theta_H|y_L)$. Thus, we cannot be in a situation in which constraint $C_{2H} \leqslant Y_{2H}$ is binding in problem (16) because at $s = 1$ we have

$$Y_1 < C_H^o(1) = Y_{2H}.$$

[24]In most numerical examples we find a cutoff ρ^* such that for $\rho < \rho^*$ it is optimal to provide no information as condition (14) holds, while for $\rho > \rho^*$ it is optimal to provide partial information and have $C_1 = C_2(y_H)$ because condition (14) does not hold. Note that we have to rely on numerical examples because in this economy ρ affects both $\pi_1(y_1)$ and the transition probabilities $\pi_2(y_2|y_1)$, while the pure exchange economy in the text allows to vary $\pi_1(y_1)$ keeping $\pi_2(y_2|y_1)$ constant.

B Omitted proofs

This appendix collects the proofs omitted from the main text, in the order in which the corresponding results appear.

B.1 Proof of Lemma 1

The IC constraint in problem 2 must be binding or the solution of the problem would be $c_{am} = Y_{2H}$ and $c_{pm} = 0$, which would clearly violate either the IC or the budget constraint. The budget constraint is binding by monotonicity of preferences. We can then substitute for C_L^o in the budget constraint using the IC constraint and further substitute

[24]Another way to see the issue is that it is not possible for $\pi_1(y_H)$ to go above $\pi_2(y_H|y_H)$ because even if $\rho \to 1$ then $\pi_1(y_H) \to \pi_H(\theta_H)$ with constant expected income in period 1 and 2.

for c_{am} in the objective function using the budget constraint. The objective function simplifies to

$$sY_{2H} + (1-s)\,Y_{2L} + [-s\pi_2(y_H|y_H) - (1-s)\,\beta\pi_2(y_H|y_L) + \beta\pi_2(y_H|y_H)]\,c_{pm}.$$

The coefficient on c_{pm} is negative iff $s \geqslant s^* \equiv \beta\kappa$. In that case, the solution to the problem entails $c_{pm} = 0$ and $c_{am} = C_L^o = sY_{2H} + (1-s)\,[Y_{2L}]$. If $s < s^*$, the objective is increasing in c_{pm}. Expressing C_L^o as a function of c_{pm}, we obtain

$$C_L^o = sY_{2H} + (1-s)\,Y_{2L} - s\,[\pi_2(y_H|y_H) - \beta\pi_2(y_H|y_L)]\,c_{pm}$$

which is strictly decreasing in c_{pm}. It follows that the participation constraint is binding, $C_L^o = Y_{2L}$, and that $c_{pm} = \kappa(y_H - y_L)$. From the IC constraint we further obtain $c_{am} = \pi_2(y_H|y_H)(y_H - y_L)(1-\kappa) + y_L$. Q.E.D.

B.2 Derivation of the equilibrium outcome with one-sided commitment

As we describe in the text, for a given disclosure policy, the optimal contract offered in period 1 solves

$$\max_{C_1,C_2} u(C_1) + \sum_{y_1}\sum_{m} \mu(m|y_1)\,\pi_1(y_1)\,u(C_2(y_1,m)) \tag{21}$$

subject to the firm's nonnegative-profit condition,

$$Y_1 - C_1 + \sum_{i=L,H}\sum_{m} \mu(m|y_i)\,\pi_1(y_i)\,(Y_{2i} - C_2(y_i,m)) \geqslant 0,$$

and the ex-post participation constraints for high- and low-income consumers (6) and (7), respectively.

The participation constraint for the low-income consumers (7) is slack because the incumbent wants to subsidize consumption of the low types in period 2. Thus, the solution to problem (21) has $C_1 = C_2(y_L, m)$ for all m, while the consumption for consumers with history (y_H, m) is given by $C_2(y_H, m) = \max\{C_1, C_H^o(s(m))\}$, and the level of C_1 is given by the nonnegative-profit condition with equality:

$$C_1(1+\pi_1(y_L)) + \pi_1(y_H)\sum_{m}\mu(m|y_H)\max\{C_1, C_H^o(s(m))\} = Y_1 + Y_2.$$

B.3 Proof of Proposition 1

Optimal disclosure policy We now turn to the design of the optimal public disclosure policy. Before proceeding, we prove Lemma 3, that $C_H^o(s)$ is convex.

Proof. For $s < s^*$, $C_H^o(s) = C_H^{LCS}$ which is a constant. For $s \geqslant s^*$, $C_H^o(s) = sY_{2H} + (1-s)Y_{2L}$ which is a linear increasing function of s. Since C_H^o is a continuous function, it immediately follows that $C_H^o(s)$ is convex. Q.E.D.

That is, the incumbent's cost of providing utility $V_H^o(s)$ in period 2, $C_H^o(s)$, is convex in the share of high-income consumers. This lemma guarantees that the maximal profits the incumbent firm makes on high-income consumers do not increase if those consumers are associated with more than one signal.

The next proposition shows that the value in (8) is maximized by a disclosure policy that provides no information.

Proof. We can equivalently write the problem of choosing the optimal disclosure policy μ as one of choosing the optimal distribution of shares s subject to the Bayesian plausibility constraint,

$$\sum_s p(s)\, s = \pi_1(y_H). \tag{22}$$

That is,

$$\max_{p \in \Delta([0,1])} \max_{C_1, V_L(s), V_H(s)} u(C_1) + \sum_s p(s)\left[sV_H(s) + (1-s)V_L(s)\right] \tag{23}$$

subject to

$$C_1 = Y_1 + Y_2 - \sum_s p(s)\left[sh(V_H(s)) + (1-s)h(V_L(s))\right] \geqslant 0$$

$$V_H(s) \geqslant V_H^o(s), \tag{24}$$

and (22) where $h = u^{-1}$. Constraint (24) is the ex-post participation constraint (6) written in utility terms.

Suppose by way of contradiction that the optimal disclosure policy induces a distribution $p \neq p^*$, where p^* is such that $p^*(\pi_1(y_H)) = 1$. If $\bar{V}_H = \sum_s p(s)\, sV_H(s)/\pi_1(y_H) \geqslant V_H^o(\pi_1(y_H))$, then by convexity of h, it is possible to deliver $\bar{V}_H$ at a lower cost with no information. Thus, by providing no information, it is possible to improve the ex-ante utility by increasing consumption in period 1 (or by increasing $V_L(s)$). If instead $\bar{V}_H = \sum_s p(s)\, sV_H(s)/\pi_1(y_H) < V_H^o(\pi_1(y_H))$, note that

$$\sum_s p(s)\, sh(V_H(s)) \geqslant \sum_s p(s)\, sh(V_H^o(s)) = \sum_s p(s)\, sC_H^o(s) \geqslant \pi_1(y_H)\, C_H^o(\pi_1(y_H)),$$

where the first inequality follows from (24), the second, from the definition of h, and the

last, from the convexity of C_H^o. Thus, providing no information and delivering $V_H = V_H^o(\pi_1(y_H))$ increase the continuation value for high-income consumers and lower its cost. Thus providing information cannot be optimal. Q.E.D.

B.4 Proof of Lemma 4

Similarly to the logic of Lemma 2, because consumers can always switch, any contract the incumbent offers in period 2 must deliver at least $C_H^o(s(m))$ to the high types and Y_{2L} to the low types. Thus, the incumbent will always offer contracts that satisfy (6) and (7). We are going to show that the following is the unique equilibrium outcome: the incumbent offers a contract with no distortion and expected consumption $C_H^o(s(m))$ to the consumers with history (y_H, m) and a contract with no distortion and $c_{2,am}(y_L, m) = C_2(y_L, m) = Y_{2L}$ to the consumers with history (y_L, m).

Suppose first that in equilibrium low-income consumers receive a payoff of $u(Y_{2L})$. It is then clear that the incumbent will offer a contract with full insurance and value $C_H^o(s(m))$ to the consumers with history (y_H, m). Fully front-loading consumption in the am (no distortions) is optimal to minimize the cost of delivering such a level of utility. Note that the incumbent is making positive profits because $C_H^o(s(m)) \leqslant Y_{2H}$, with strict inequality unless the signal is fully revealing. Offering $C_H^o(s(m))$ is optimal because if the incumbent offers less, then the outsiders can attract all high-income consumers and erase all the incumbent's profits. Offering more is not optimal, because it only reduces profits.

We are left to show that it is optimal to offer a contract with no distortions with value $u(Y_{2L})$ to the low-income consumers. Clearly, offering less is not feasible, because any outsider can always offer a contract with no distortions and value $u(Y_{2L})$. Offering more has no advantages, because the outside options for the high-income consumers are constant in the value offered by the incumbent to the low-income consumers. Q.E.D.

B.5 Derivation of the equilibrium outcome with two-sided lack of commitment

As we show in Lemma 4 and 5, given a public disclosure policy (M, μ), the equilibrium outcome is given by

$$C_1 = Y_1 + \pi_1(y_H) \sum_m \mu(m|y_H) \Pi(m) \tag{25}$$

$$C_2(y_L, m) = Y_{2L} \tag{26}$$

$$C_2(y_H, m) = C_H^o(s(m)), \tag{27}$$

where $\Pi(m) = Y_{2H} - C_H^o(s(m))$.

As discussed in the remark in Section 2, if the withdrawal cost is strictly positive and the informed incumbent acts as a Stackelberg leader, this is the unique equilibrium outcome. To understand the role played by the Stackelberg leader assumption, suppose that $s(m) > s^*$ so that $C_L^o(s(m)) > Y_{2L}$. When the informed incumbent is a Stackelberg leader, it offers $C_H^o(s(m))$ to high types to prevent outsiders from offering a menu $\{C_H^o(s(m)), C_L^o(s(m))\}$ and poaching high-income consumers.

Consider instead the case where all firms act simultaneously. If the menu $\{C_H^o(s(m)), C_L^o(s(m))\}$ is offered by at least one outsider, the informed incumbent then offers $C_H^o(s(m))$ to the high types to retain them, since it makes strictly positive profits on them because of its informational advantage. If the withdrawal cost is strictly positive, this cannot be an equilibrium: offering $\{C_H^o(s(m)), C_L^o(s(m))\}$ is not optimal for an outsider if the informed incumbent offers $C_H^o(s(m))$ or more to the high types. This is because all high types stay with the incumbent, while all the low types accept the outsider's offer, if not withdrawn. Thus, this strategy earns the outsider strictly negative profits,

$$\max\{-\varepsilon, -\pi_1(y_L)\mu(m|y_L)(C_L^o(s(m)) - Y_{2L})\} < 0.$$

Offering nothing is then preferable for the outsiders. If the incumbent offers less than $C_H^o(s(m))$ to the high types because it expects the outsiders to offer less than that, then it is optimal for the outsiders to offer $\{C_H^o(s(m)), C_L^o(s(m))\}$. Then no pure strategy equilibrium exists.

If the withdrawal costs ε are equal to zero, the outcome in Lemma 5 is also an equilibrium outcome when all firms act simultaneously. To support this equilibrium outcome, outsiders offer the menu $\{C_H^o(s(m)), C_L^o(s(m))\}$ and withdraw their offer if the informed incumbent offers $C_H^o(s(m))$ to the high types. The outsiders then attains a payoff of 0,

so they are indifferent between this choice and not offering anything. However, if $\varepsilon = 0$, there can be multiple equilibria even absent the informed incumbent, as shown in Netzer and Scheuer (2014). The equilibrium outcome we consider is the one that minimizes the incumbent's period 2 profits.

B.6 Proof of Proposition 2

First we show that there exists a π^* such that condition (14) holds if and only if $\pi_1(y_H) \leqslant \pi^*$. To simplify the notation, we write $\pi_1 = \pi_1(y_H)$ when not ambiguous. Letting

$$K(s) \equiv C_H^o(s),$$

we can re-write condition (14) as

$$f(\pi_1) \equiv K(\pi_1) - [\pi_1 y_H + (1-\pi_1) y_L + \pi_1 (Y_{2H} - K(\pi_1))] \geqslant 0$$

Note that f is convex because K is increasing and convex, so $(1+\pi_1) K(\pi_1)$ is convex in π_1 and all other terms are affine. Furthermore, evaluating at $\pi_1 = 0$ and $\pi_1 = \pi_2(y_H|y_H)$ we have

$$\begin{aligned} f(0) &= K(0) - y_L > 0 \\ f(\pi_2(y_H|y_H)) &= -(1+\pi_2(y_H|y_H))(Y_{2H} - K(\pi_2(y_H|y_H))) < 0. \end{aligned}$$

Thus, since $f(0) > 0$, $f(\pi_2(y_H|y_H)) < 0$ and f is convex, f must cross 0 only once between $(0, \pi_2(y_H|y_H))$, and we denote such point by π^*. Then for $\pi_1 \in (0, \pi^*)$ we have $f(\pi_1) > 0$ and for $\pi \in (\pi^*, \pi_2(y_H|y_H))$ we have $f(\pi_1) < 0$.

Suppose that $\pi_1(y_H) \geqslant \pi^*$ so condition (14) does not hold. First, we show that all high-income consumers receive the same signal. Suppose by way of contradiction that the optimal disclosure policy, (M^*, μ^*), is such that there are two signals, m_1 and m_2, with $\mu^*(m_1|y_H) > 0$ and $\mu^*(m_2|y_H) > 0$ and $V_H^o(s(m_1)) \neq V_H^o(s(m_2))$. Let

$$\bar{V} = \sum_m \mu^*(m|y_H) V_H^o(s(m)).$$

Suppose that $\bar{V} \geqslant V_H^o(\pi_1(y_H))$. In this case, there exists an alternative disclosure policy with $M = \{g, b\}$ with $\mu(g|y_H) = 1$ and $\mu(b|y_L) \in [0,1)$ and $V_H^o(s(g)) = \bar{V}$. Due to concavity of the utility function, this alternative disclosure allows the incumbent to save resources in period 2 and still deliver expected utility $\bar{V}$ to high-income consumers.

Because of competition in period 1, these additional profits are rebated to the consumer in the first period. Thus, this alternative disclosure policy is an improvement because it delivers the same expected utility in period 2 but higher utility in period 1, reaching a contradiction.

Consider now the case with $\bar{V} < V_H^o(\pi_1(y_H))$. Consider an alternative policy that discloses no information. The value in the second period for the high-income type is higher than $\bar{V}$. Also utility in period 1 is higher. In fact, note that

$$\begin{aligned}
&u\left(Y_1 + \pi_1(y_H)\sum_{m\in M}\mu(m|y_H)\left[Y_{2H} - h\left(V_H^o(s(m))\right)\right]\right) + \pi_1(y_H)\bar{V} \\
<&u\left(Y_1 + \pi_1(y_H)\left[Y_{2H} - h\left(\bar{V}\right)\right]\right) + \pi_1(y_H)\bar{V} \\
<&u\left(Y_1 + \pi_1(y_H)\left[Y_{2H} - h\left(V_H^o(\pi_1(y_H))\right)\right]\right) + \pi_1(y_H)V_H^o(\pi_1(y_H))
\end{aligned}$$

where the first inequality follows from the convexity of h, and the second inequality from the concavity of u, convexity of h and condition (14) not holding. To see this last step, note that

$$F(V) \equiv u\left(Y_1 + \pi_1(y_H)\left[Y_{2H} - h(V)\right]\right) + \pi_1(y_H)V$$

is increasing in V for $V < V_H^o(\pi_1(y_H))$ when (14) fails. In fact,

$$\begin{aligned}
F'(V) &\equiv \pi_1(y_H)\left[1 - u'\left(Y_1 + \pi_1(y_H)\left[Y_{2H} - h(V)\right]\right)h'(V)\right] \\
&> \pi_1(y_H)\left[1 - u'\left(Y_1 + \pi_1(y_H)\left[Y_{2H} - h\left(V_H^o(\pi_1)\right)\right]\right)h'\left(V_H^o(\pi_1)\right)\right] \\
&= \pi_1(y_H)\left[1 - \frac{u'\left(Y_1 + \pi_1(y_H)\left[Y_{2H} - h\left(V_H^o(\pi_1)\right)\right]\right)}{u'\left(h\left(V_H^o(\pi_1)\right)\right)}\right] \\
&> 0
\end{aligned}$$

where the first inequality follows from u being concave, h increasing and convex and $V_H^o(\pi_1) > V$; and the last inequality from condition (14) not holding that implies

$$u'\left(Y_1 + \pi_1(y_H)\left[Y_{2H} - h\left(V_H^o(\pi_1)\right)\right]\right) < u'\left(h\left(V_H^o(\pi_1)\right)\right).$$

Thus, even in this case the alternative disclosure policy improves upon the original allocation yielding a contradiction.

We established that when condition (14) fails all consumers with high income in period 1 receive the same signal. Without loss of generality we can consider $M = \{g, b\}$ and have $\mu(g|y_H) = 1$. We are only left to choose the fraction of consumers with low income in period 1 that receive the same "good" signal. This is not going to affect their consump-

tion but it affects the composition of the pool of agents with a good signal and therefore the continuation value of the high income. The set of implementable continuation values for the high income consumer is $[V_H^o(\pi_1(y_H)), V_H^o(1)]$. We can then write the problem as

$$\max_{c_1, V_H} u(c_1) + \pi_1(y_H) V_H + \pi_1(y_L) u(Y_{2L}) \tag{28}$$

subject to

$$Y_1 - c_1 + \pi_1(y_H)[Y_{2H} - h(V_H)] \geqslant 0$$

$$V_H \in [V(\pi_1(y_H)), V_H^o(1)]$$

and then recover $\mu(b|y_L)$ from

$$V_H = V_H^o\left(\frac{\pi_1(y_H)}{\pi_1(y_H) + \pi_1(y_L)(1 - \mu(b|y_L))}\right). \tag{29}$$

If the last constraint in problem (28) does not bind and the optimal V_H is interior, it is clear that $u'(c_1) = u'(c_{2H})$ and so

$$C_1 = C_{2H} = \bar{C} \equiv \frac{Y_1 + \pi_1(y_H) Y_{2H}}{1 + \pi_1(y_H)} > Y_1$$

Thus, $\mu(b|y_L)$ solves

$$V_H^o\left(\frac{\pi_1(y_H)}{\pi_1(y_H) + \pi_1(y_L)(1 - \mu(b|y_L))}\right) = u(\bar{C})$$

Then we just have to check if $u(\bar{C}) \in [V_H^o(\pi_1(y_H)), V_H^o(1)]$. Because condition (14) is not satisfied, $V_H^o(\pi_1(y_H)) \leqslant u(\bar{C})$. Moreover, since $\pi_1 \leqslant \pi_2(y_H|y_H)$ then $u(\bar{C}) \leqslant V_H^o(1)$. Thus, the last constraint does not bind.

Finally, suppose that $\pi_1(y_H) < \pi^*$ so condition (14) is satisfied. In this case, it is not possible to equalize consumption in period 1 and period 2 if $y_1 = y_H$ by assigning the same signal to all high-income consumers. This is because the last constraint in (28) binds and $u'(C_1) > u'(C_2(y_H))$. It might then be optimal to assign different signals to the high-income consumers in order to reduce their expected continuation value, economize on resources used in period 2, increase the profits rebated in period 1. This is not feasible since $K(s) = h(V(s))$ is convex. Suppose by way of contradiction that the optimal disclosure policy, (M^*, μ^*), is such that there are two signals, m_1 and m_2, with $\mu^*(m_1|y_H) > 0$

and $V_H^o(s(m_1)) \neq V_H^o(s(m_2))$ with

$$\bar{V} = \sum_m \mu^*(m|y_H) V_H^o(s(m)) < V_H^o(\pi_1(y_H)).$$

(Clearly, if $\bar{V} > V_H^o(\pi_1)$ an argument similar to the one above shows that assigning multiple signals to high income consumers is not optimal.) The period 1 consumption associated with this plan is

$$\begin{aligned}
&Y_1 + \pi_1(y_H)\left(Y_{2H} - \sum_m \mu^*(m|y_H) K(s(m))\right) \\
\leqslant &Y_1 + \pi_1(y_H)\left(Y_{2H} - K\left(\sum_m \mu^*(m|y_H) s(m)\right)\right) \\
\leqslant &Y_1 + \pi_1(y_H)(Y_{2H} - K(\pi_1(y_H)))
\end{aligned}$$

where the first inequality follows from the convexity of K, and the second from the observation that $\sum_m \mu^*(m|y_H) s(m) \geqslant \pi_1(y_H)$ and K is increasing. Thus, disclosing no information increases both the expected continuation value in period 2 for the high-income consumers and consumption in period 1. Thus, the original allocation cannot be optimal and it must be optimal to assign the same signal to all high-income consumers. Therefore, problem (28) characterizes the full problem. Since the last constraint is binding, one way to obtain the optimum is to provide no information so $C_2(y_H) = h(V_H^o(\pi_1))$ and

$$C_1 = Y_1 + \pi_1(Y_{2H} - h(V_H^o(\pi_1))) < C_2(y_H).$$

Q.E.D.

B.7 Proof of Lemma 6

Problem 19 can be solved using the graphic approach in Kamenica and Gentzkow (2011). The value of the objective function is the convex hull of $W(s)$ evaluated at $\pi_1(y_H)$. We report the graph of $W(s)$ in Figure 7 for convenience. The key step is showing that the convex hull of $W(s)$ is equal to the line connecting $W(0)$ to $W(s^*)$ for $s < s^*$ and it is equal to $W(s)$ for $s \geqslant s^*$. It then follows that for $s < s^*$ the optimal information disclosure entails two signals with composition $s = 0$ and $s = s^*$, while for $s \geqslant s^*$ the optimal disclosure has a single pool of quality $\pi_1(y_H)$, which concludes the proof.

Since $W(s)$ is strictly concave for $s > s^*$, we need to show that the right derivative of $W(s)$ at s^* is smaller than the slope of the line connecting $W(0)$ to $W(s^*)$. The former is

Figure 7: Expected continuation value with spot contracts, $W(s)$

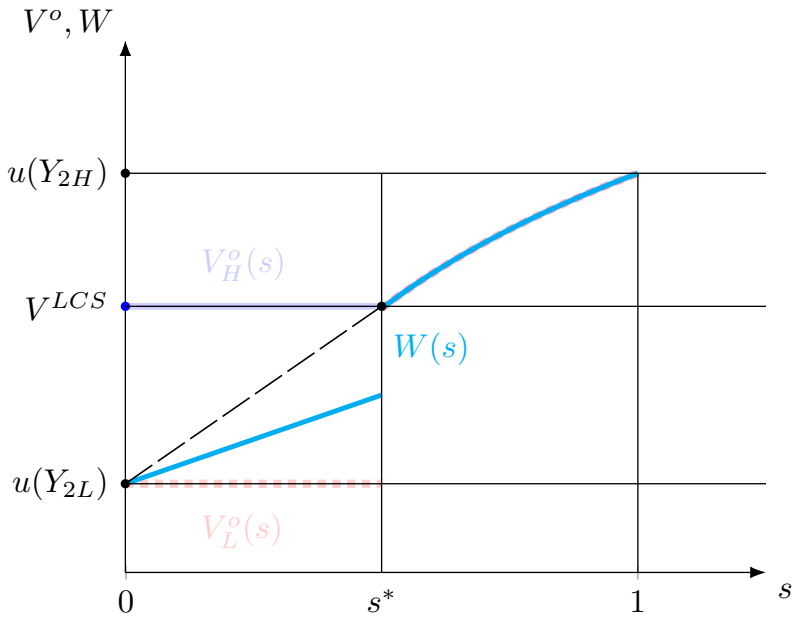

equal to $u'(Y(s^*))(Y_{2H} - Y_{2L})$, while the latter is equal to

$$\frac{u(Y(s^*)) - u(Y_{2L})}{s^*}$$

where $Y(s) \equiv sY_{2H} + (1-s)Y_{2L}$. Strict concavity of u then implies

$$\frac{u(Y(s^*)) - u(Y_{2L})}{Y(s^*) - Y_{2L}} > u'(Y(s^*)) \Rightarrow \frac{u(Y(s^*)) - u(Y_{2L})}{s^*} > u'(Y(s^*))(Y_{2H} - Y_{2L})$$

Q.E.D.

B.8 Proof of Lemma 7

To characterize the optimum, it is convenient to first define the expected continuation value for consumers with a signal characterized by the share s:

$$W(s;\Delta) = sV_H^o(s) + (1-s)V_L(s) + \Delta,$$

where

$$V_L(s) = \begin{cases} V_L^o(s) & \text{if } s \geqslant \hat{s}(\Delta) \\ V_L^o(0) = u(Y_{2L}) & \text{if } s < \hat{s}(\Delta) \end{cases}.$$

The optimal public disclosure policy solves

$$\max_{C_1, p(s) \in \Delta([0,1])} u(C_1) + \sum_s p(s) W(s;\Delta)$$

subject to

$$\sum_s p(s)\, s = \pi_1$$

$$C_1 = Y_1 + \sum_s p(s)\, s \left[Y_{2H} - \min\left\{h\left(V_H^o(s,\Delta)\right), Y_{2H}\right\}\right]$$

Suppose by way of contradiction that p^* that solves the problem above assigns positive probabilities to more than one shares above $\hat{s}(\Delta)$. That is, there are $\{s_1, s_2, ..., s_N\}$ with $N \geqslant 2$ such that $p^*(s_n) > 0$ and $s_n \geqslant \hat{s}(\Delta)$ for all $n = 1, ..., N$. Consider an alternative solution that only assigns probability $p(\bar{s}) = \sum_{n=1}^N p^*(s_n)$ to $\bar{s} = \sum_{n=1}^N \frac{p^*(s_n)}{p(\bar{s})} s_n$ for $s \geqslant \hat{s}(\Delta)$ and $p(s) = p^*(s)$ for all $s < \hat{s}(\Delta)$. Since the incumbent does not make any profits on consumers with $s \geqslant \hat{s}(\Delta)$ then this alternative solution does not change C_1. However, by concavity of $W(s;\Delta)$ for $s \geqslant \hat{s}(\Delta)$ we have that

$$\sum_{n=1}^N p^*(s_n) W(s_n;\Delta) = p(\bar{s}) \sum_{n=1}^N \frac{p^*(s_n)}{p(\bar{s})} W(s_n;\Delta)$$

$$< p(\bar{s})\, W\left(\sum_{n=1}^N \frac{p^*(s_n)}{p(\bar{s})} s_n; \Delta\right) = p(\bar{s})\, W(\bar{s};\Delta)$$

so the alternative solution attains a higher value. Hence we have a contradiction.

The logic we used in Proposition 2 to show that all good types that stay have the same signal can be replicated to show that there is at most one group with $s \in (0, \hat{s}(\Delta))$ that are retained by the incumbent. Q.E.D.

B.9 Proof of Proposition 3

Given Lemma 7, we can then write the problem for the optimal information disclosure policy as

$$\max_{s_{in}, s_{out}, p_0, p_{in}, p_{out}} u\left(Y_1 + p_{in} s_{in}\left[Y_{2H} - h\left(V_H^o(s_{in}) + \Delta\right)\right]\right)$$

$$+ p_0 W(0,\Delta) + p_{in} W(s_{in},\Delta) + p_{out} W(s_{out},\Delta)$$

subject to

$$\begin{aligned} p_0 + p_{in} + p_{out} &= 1, \quad p_s \geqslant 0 \\ p_{in} s_{in} + p_{out} s_{out} &= \pi_1 (y_H) \\ s_{in} \in (0, \hat{s}(\Delta)), &\quad s_{out} \in [\hat{s}(\Delta), 1] \end{aligned}$$

where recall that $h = u^{-1}$.

First we show that there exists a $\Delta_L > 0$ such that $p_{out} = 0$ for $\Delta \leqslant \Delta_L$. Let Δ be close to 0. Suppose by way of contradiction that $p_{out} > 0$. Let V be the value at the optimum

$$\begin{aligned} V = u(Y_1 + p_{in} s_{in} \Pi(s_{in})) + p_{in} [s_{in} V_H^o(s_{in}) + (1 - s_{in}) V_L^o(0)] \\ + p_{out} [s_{out} V_H^o(s_{out}) + (1 - s_{out}) V_L^o(s_{out})] + p_0 W(0) + \Delta. \end{aligned}$$

Consider a deviation to an alternative policy $(s'_{in}, p'_0, p'_{in}, p'_{out} = 0)$ that delivers the same expected value to the high-income consumers as the one with $p_{out} > 0$, $\sum_{i=in,out} p_i s_i V_H(s_i) = \pi_1 V_H(s'_{in})$. For Δ close to 0, such policy exists and $s'_{in} < \hat{s}(\Delta)$. Then notice that, by (strict) convexity of $K(s)$,

$$\sum_{i=in,out} p_i s_i h(V_H^o(s_i) + \Delta) > p'_{in} s'_{in} h(V_H^o(s'_{in}) + \Delta)$$

As $\Delta \to 0$ we have that $s_{out} \to 1$ and $h(V_H(1)) = Y_{2H}$, so the expected profits in the original allocation are

$$\begin{aligned} p_{in} s_{in} \Pi(s_{in}) &= p_{in} s_{in} Y_{2H} - p_{in} s_{in} h(V_H^o(s_{in})) \\ &= \pi_1 Y_{2H} - \sum_{i=in,out} p_i s_i h(V_H^o(s_i)) \\ &< \pi_1 Y_{2H} - p'_{in} s'_{in} h(V_H^o(s'_{in})) \\ &= p'_{in} s'_{in} \Pi(s'_{in}) \end{aligned}$$

where the second equality follows from $p_{out} s_{out} (Y_{2H} - h(V_H(1))) = 0$. Then,

$$V' - V = u(Y_1 + p'_{in} s'_{in} \Pi'(s'_{in})) - u(Y_1 + p_{in} s_{in} \Pi(s_{in})) > 0$$

since, by construction, the continuation value of the high-income consumers is equal across policies, and all low-income consumers receive $u(Y_{2L})$ under both policies. As the deviation is profitable, we have reached a contradiction. Thus, for Δ close enough to

0 it is optimal to have $p_{out} = 0$. The intuition is that when Δ is small enough, a positive p_{out} is akin to providing full information about a subset of high types in our baseline model, which we established not to be optimal.

Next, we show that there exist $\Delta_H < \bar{\Delta}$ such that $p_{in} = 0$ for $\Delta \geqslant \Delta_H$. This is the case since $p_{in} = 0$ for $\Delta \geqslant \bar{\Delta}$. By definition of $\bar{\Delta}$, $V_H^{LCS} + \bar{\Delta} = u(Y_{2H})$ so for all $s \leqslant \hat{s}(\bar{\Delta})$ the incumbent makes no profits on the high-income consumers it retains. Then $C_1 = Y_1$. We want to show that if Δ is sufficiently close to $\bar{\Delta}$ then it is optimal to have $p_{in} = 0$. Suppose by way of contradiction that $p_{in} > 0$. The value is

$$V = u(Y_1 + p_{in} s_{in} \Pi(s_{in})) + p_{in} W(s_{in}) + p_{out} W(s_{out}) + p_0 W(0)$$

Consider a perturbation that decreases p_{in} by $\varepsilon > 0$, increases p_{out} by ε, and decreases s_{out} by $\varepsilon(s_{out} - s_{in})/p_{out}$ so that the perturbation is feasible. The marginal value of such perturbation is

$$\begin{aligned} dV = &-s_{in} \Pi(s_{in}) u'(Y_1 + p_{in} s_{in} \Pi(s_{in})) - W(s_{in}) \\ &+ W(s_{out}) - W'(s_{out})(s_{out} - s_{in}) \end{aligned}$$

As $\Delta \to \bar{\Delta}$ we have $\Pi(s_{in}) \to 0$ so

$$\begin{aligned} dV &\to W(s_{out}) - W(s_{in}) - W'(s_{out})(s_{out} - s_{in}) \\ &\geqslant W(s_{out}) - \bar{W}(s_{in}) - W'(s_{out})(s_{out} - s_{in}) > 0 \end{aligned}$$

The first inequality follows from $W(s_{in}) \leqslant \bar{W}(s_{in})$ where $\bar{W}$ is the concavification of $W(s)$. The second inequality follows from $s_{out} \geqslant \hat{s}(\Delta) \geqslant s^*$, $s_{out} \geqslant \hat{s}(\Delta) > s_{in}$, and strict concavity of W over $[s^*, 1]$ as graphically illustrated in Figure 4. Thus, the perturbation is welfare improving for Δ sufficiently close to $\bar{\Delta}$, reaching a contradiction. Then it must be that $p_{in} = 0$ for all Δ close to $\bar{\Delta}$.

Finally, to see that full information is never optimal, simply note that a feasible policy solves the static problem (19) with the additional constraint $s_{out} \geqslant \hat{s}(\Delta)$, namely $s_{out} = \max\{\hat{s}(\Delta), \pi_1\}$, $p_{out} = \pi_1 / s_{out}$, $s_{in} = 0$. The value associated with this feasible solution is

$$\begin{aligned} &u(Y_1) + [(1 - p_{out}) W(0, \Delta) + p_{out} W(s_{out}, \Delta)] \\ >&u(Y_1) + [(1 - \pi_1) W(0, \Delta) + \pi_1 W(1, \Delta)] \end{aligned}$$

since W is strictly concave over $[\hat{s}(\Delta), 1]$ and $s_{out} < 1$ whenever $\pi_1 < 1$. Hence this

alternative feasible solution attains a higher value than full information disclosure. Thus, full information is never optimal. Q.E.D.

B.10 Proof of Proposition 4

To simplify notation, let $\pi_L \equiv \pi_2(y_H|y_L) < \pi_H \equiv \pi_2(y_H|y_H)$. Thus, $\pi_1 \equiv \pi_1(y_H) \in [\pi_L, \pi_H]$. Consider first $\pi_1 = \pi_L$. In this case, $Y_1 = Y_{2L}$, and by Proposition 2 the maximal value under long-term relations is no disclosure. Consider two cases. First, if $\pi_L < s^*$ then with long-term relations, $C_{2H} = C_H^{LCS} = s^* Y_{2H} + (1 - s^*) Y_{2L}$ and

$$C_1 = Y_{2L} + \pi_L (Y_{2H} - C_{2H}) = Y_{2L} + \pi_L (1 - s^*)(Y_{2H} - Y_{2L})$$

so

$$V^{LT}(\pi_L) = u(C_1) + \pi_L u\left(C_H^{LCS}\right) + (1 - \pi_L) u(Y_{2L})$$

With short-term contracts under the optimal disclosure policy[25]

$$V^{ST}(\pi_L) = u(Y_{2L}) + \frac{\pi_L}{s^*} u(s^* Y_{2H} + (1 - s^*) Y_{2L}) + \left(1 - \frac{\pi_L}{s^*}\right) u(Y_{2L})$$

Then

$$V^{LT}(\pi_L) - V^{ST}(\pi_L) = u(C_1) - u(Y_{2L}) - \frac{\pi_L (1 - s^*)}{s^*} \left[u\left(C_H^{LCS}\right) - u(Y_{2L})\right]$$

Note that, defining $\lambda = \frac{\pi_L (1 - s^*)}{s^*} \in (0, 1)$, we can express

$$C_1 = (1 - \lambda) Y_{2L} + \lambda C_H^{LCS}$$

so by strict concavity of u:

$$\begin{aligned} V^{LT}(\pi_L) - V^{ST}(\pi_L) &= u\left((1 - \lambda) Y_{2L} + \lambda C_H^{LCS}\right) - \left[\lambda u\left(C_H^{LCS}\right) + (1 - \lambda) u(Y_{2L})\right] \\ &> u\left((1 - \lambda) Y_{2L} + \lambda C_H^{LCS}\right) - \left[u\left(\lambda C_H^{LCS} + (1 - \lambda) Y_{2L}\right)\right] = 0. \end{aligned}$$

Second, if $\pi_L \geqslant s^*$ then with long-term relations, $C_{2H} = \pi_L Y_{2H} + (1 - \pi_L) Y_{2L}$ and

$$C_1 = Y_{2L} - \pi_L (Y_{2H} - C_{2H}) = Y_{2L} + \pi_L (1 - \pi_L)(Y_{2H} - Y_{2L})$$

[25] Note that at the optimum the share of high type with a good signal is $s^* = \pi_H / [\pi_H + (1 - \pi_H) \mu(g|y_L)]$. Thus, the mass of low type with a bad signal is $(1 - \pi_H)(1 - \mu(g|y_L)) = 1 - \pi_H / s^*$.

and, since $C_{2H} - Y_{2L} = \pi_L (Y_{2H} - Y_{2L})$, then

$$C_1 = \pi_L Y_{2L} + (1 - \pi_L) C_{2H}$$

and

$$V^{LT}(\pi_L) = u(C_1) + \pi_L u(C_{2H}) + (1 - \pi_L) u(Y_{2L})$$

With short-term contracts under the optimal disclosure policy

$$V^{ST}(\pi_L) = u(Y_{2L}) + u(\pi_L Y_{2H} + (1 - \pi_L) Y_{2L}) = u(Y_{2L}) + u(C_{2H})$$

Then, by strict concavity of u:

$$\begin{aligned} V^{LT}(\pi_L) - V^{ST}(\pi_L) &= u(\pi_L Y_{2L} + (1 - \pi_L) C_{2H}) + \pi_L u(C_{2H}) + (1 - \pi_L) u(Y_{2L}) \\ &\quad - [u(Y_{2L}) + u(C_{2H})] \\ &> \pi_L u(Y_{2L}) + (1 - \pi_L) u(C_{2H}) + \pi_L u(C_{2H}) + (1 - \pi_L) u(Y_{2L}) \\ &\quad - [u(Y_{2L}) + u(C_{2H})] = 0 \end{aligned}$$

Hence, in both cases, $V^{LT}(\pi_L) > V^{ST}(\pi_L)$. By continuity, if π_1 is sufficiently close to π_L then

$$V^{LT}(\pi_1) > V^{ST}(\pi_1).$$

Consider next $\pi_1 = \pi_H$. In this case, $Y_1 = Y_{2H}$ and by Proposition 2 the maximal value under long-term relations is

$$V^{LT}(\pi_H) = u(Y_{2H}) + \pi_H u(Y_{2H}) + (1 - \pi_H) u(Y_{2L})$$

with perfect smoothing between period 1 and period 2 after a good income realization. Thus, $V^{LT}(\pi_H)$ is the same as the full information outcome with static contracts. Consider now the problem under repeated static contracts. There are two cases. First, if $\pi_H \geqslant s^*$ then by Lemma 6 there is no disclosure and

$$\begin{aligned} V^{ST}(\pi_H) &= u(Y_{2H}) + u(\pi_H Y_{2H} + (1 - \pi_H) Y_{2L}) \\ &> u(Y_{2H}) + \pi_H u(Y_{2H}) + (1 - \pi_H) u(Y_{2L}) = V^{LT}(\pi_H) \end{aligned}$$

by strict concavity of u. Second, if $\pi_H < s^*$ then by Lemma 6

$$\begin{aligned} V^{ST}(\pi_H) &= u(Y_{2H}) + \frac{\pi_H}{s^*} u(s^* Y_{2H} + (1-s^*) Y_{2L}) + \left(1 - \frac{\pi_H}{s^*}\right) u(Y_{2L}) \\ &> u(Y_{2H}) + \pi_H u(Y_{2H}) + (1-\pi_H) u(Y_{2L}) = V^{LT}(\pi_H) \end{aligned}$$

by strict concavity of u. Hence, in both cases, $V^{ST}(\pi_H) > V^{LT}(\pi_H)$. By continuity, if π_1 is sufficiently close to π_H then

$$V^{ST}(\pi_1) > V^{LT}(\pi_1).$$

Q.E.D.

C Observed outcome and unobserved signals

Throughout this appendix, $h = u^{-1}$ denotes the inverse of the utility function.

In the main text we assume that y_1 is not observable by the outsiders. Here we sketch how our analysis extends to the case in which y_1 is observable to outsiders but the incumbent and the consumer receive a signal $\tilde{y}_1$ in period 1 that is informative about the distribution of income in period 2. Assume that $\tilde{y}_1 \in \{\tilde{y}_H, \tilde{y}_L\}$ and let $p(\tilde{y}_1|y_1)$ be the probability of getting signal $\tilde{y}_1$ given y_1. Further, let $\pi_2(y_2|y_1, \tilde{y}_1)$ be the probability of y_2 given $(y_1, \tilde{y}_1)$.[26]

The outside option for the consumer with a high-signal and observable income in period 1 y_z is

$$C_z^o(s) = \max_{c_{am}, c_{pm}, C_{Lz}^o} c_{am} + \beta \pi_2(y_2|y_z, \tilde{y}_H) c_{pm}$$

subject to

$$\begin{aligned} s[Y_{2zH} - (c_{am} + \pi_2(y_2|y_z, \tilde{y}_H) c_{pm})] + (1-s)[Y_{2zL} - C_{Lz}^o] &\geqslant 0 \\ c_{am} + \beta \pi_2(y_2|y_z, \tilde{y}_L) c_{pm} &\leqslant C_{Lz}^o \\ C_{Lz}^o &\geqslant Y_{2zL} \end{aligned}$$

[26] Appendix E offers another example where the incumbent and the consumer/worker see a private signal about the return on the private investment in general human capital.

where

$$Y_{2zH} \equiv \sum_{y_2} \pi_2 (y_2|y_z, \tilde{y}_H)\, y_2$$
$$Y_{2zL} \equiv \sum_{y_2} \pi_2 (y_2|y_z, \tilde{y}_L)\, y_2.$$

It is straightforward to show that under two-sided lack of commitment for an arbitrary (M, μ) we have:

Lemma 8. *Given a public disclosure policy* (M, μ)*, the equilibrium outcome has*

$$C_1 = Y_1 + \sum_z \pi_1 (y_z)\, p (\tilde{y}_H|y_z) \sum_m \mu (m|y_z, \tilde{y}_H) [Y_{2zH} - C_z^o (s (m))] \tag{30}$$
$$C_2 (y_z, \tilde{y}_L, m) = Y_{2zL} \tag{31}$$
$$C_2 (y_z, \tilde{y}_H, m) = C_z^o (s (m)) \tag{32}$$

where $C_1 = c_{1,am}$*,* $C_2 (y_z, \tilde{y}_L, m) = c_{2,am} (y_z, \tilde{y}_L, m)$ *and* $c_{t,pm}$ *is always 0.*

Thus, the optimal disclosure policy can be found as the solution to the analog of problem (16):

$$\max_{C_1, C_{HH}, C_{LH}} u (C_1) + \pi_1 (y_H) [p (\tilde{y}_H|y_H)\, u (C_{HH}) + p (\tilde{y}_L|y_H)\, u (Y_{2HL})] + \pi_1 (y_L) [p (\tilde{y}_H|y_L)\, u (C_{LH}) + p (\tilde{y}_L|y_L)\, u (Y_{2LL})]$$

subject to

$$C_1 = Y_1 + \sum_{z=H,L} \pi_1 (y_z)\, p (\tilde{y}_H|y_z) [Y_{2zH} - C_{zH}]$$
$$C_{zH} \in [C_z^o (p (\tilde{y}_H|y_z)), Y_{2zH}], \quad z = H, L$$

and then recover $s_z (g)$ and $\mu (b|y_z, \tilde{y}_L)$ from

$$C_{zH} = C_z^o (s_z (g)) \text{ and } s_z (g) = \frac{p (\tilde{y}_H|y_z)}{p (\tilde{y}_H|y_z) + p (\tilde{y}_L|y_z) (1 - \mu (b|y_z, \tilde{y}_L))}. \tag{33}$$

D Discrimination across consumers

Throughout this appendix, $h = u^{-1}$ denotes the inverse of the utility function.

So far we have assumed that the incumbent firm cannot discriminate among consumers with the same history. We now relax this assumption and show that the optimal signal structure maintains the same features as in the restricted case. Our main interest is in the differences in the contract structure induced by discrimination. We show that when releasing information is optimal, the incumbent firm provides $C_2(y_H) < C_H^o(s(m))$ to almost all consumers with $y_1 = y_H$.

First consider the case $V_H^o(s(m)) = V_H^{LCS}$. Then the allocation coincides with that in Lemma 5,

$$C_2(y_H, m) = C_H^{LCS},$$

and lack of discrimination has no impact on the allocation. The second and more interesting case is one in which $V_H^o(s(m)) > V_H^{LCS}$. In this case, discrimination allows the incumbent firm to offer less than $V_H^o(s(m))$ to almost all high types by exploiting the fact that lower offers from outsiders attract a worse pool of agents. Recall that s^* is the largest share of high income consumers s such that

$$V_H^o(s) = V_H^{LCS}.$$

We then order all high income consumers and index them by $i \in [0,1]$. We define $i^*(m)$ to be the measure of high types with signal m that generates the share s^*, for given information policy $\mu(m|y_L)$. The value $i^*(m) \in (0,1)$ identifies two groups of consumers. For $i \in [0, i^*(m)]$, the value consumer i receives is

$$V(i) = V_H^{LCS}.$$

This is because the least cost separating contract constitutes a lower bound for all high income consumers. For $i \in [i^*(m), 1]$,

$$V(i) = V_H^o(s_i(m))$$

where

$$s_i(m) = \frac{\pi_1(y_H)\int_0^i d\tilde{i}}{\pi_1(y_H)\int_0^i d\tilde{i} + (1-\pi_1(y_H))\,\mu(m|y_L)}$$

is the share of high income consumers in a pool that includes all low types with signal m and all high types with signal m and index smaller than i. Next we solve for the optimal disclosure policy.

Optimal public disclosure

The introduction of discrimination modifies the allocation associated to a given signal structure in two ways. First, as long as the value of the Miyazaki-Wilson contract is higher than V_H^{LCS}, high-type consumers with the same signal receive unequal consumption in the second period. In particular, their utility in the second period belongs to the interval $\left[V_H^{LCS}, V_H^o\left(s\left(m\right)\right)\right]$. Second, firms are able to extract additional profits in the second period and rebate them to the agent in the first period. We show that these two features do not fundamentally change the nature of the optimal disclosure policy.

Proposition 5. *The optimal disclosure policy has a bad-news structure i.e.* $M = \{g, b\}$ *(good or bad) and* $\mu\left(g|y_H\right) = 1$ *and* $\mu\left(b|y_L\right) \in [0,1]$*. i) For* π_1 *sufficiently low, it is optimal to provide no information; ii) For all* π_1*, more information is disclosed under discrimination–and strictly so if* π_1 *is sufficiently high.*

Proof. Proposition 5 shares most predictions with Proposition 2. Part i) follows from $y_L < h\left(V_H^{LCS}\right)$ since $h\left(V_H^{LCS}\right)$ is the minimum consumption that must be guaranteed to high-income consumers. Part ii) hinges on the fact that the incumbent firm earns more profits in the second period under discrimination, hence consumption is typically more front-loaded for all interior disclosure policies.

To ease notation, we replace $\mu\left(b|y_L\right)$ with μ, and, with slight abuse of notation, we replace the dependence of $i\left(.\right)$ and $s_i\left(.\right)$ on m with that on the associated μ. Formally, the optimal signal under discrimination solves

$$\max_{\mu\in[0,1]} u\left(Y_1 + \pi_1\left(y_H\right)\Pi\left(\mu\right)\right) + \pi_1\left(y_H\right)\left[i^*\left(\mu\right)V_H^{LCS}\left(y_H\right) + \int_{i^*(\mu)}^1 V_H^o\left(s_i\left(\mu\right)\right)di\right]$$

where

$$\Pi\left(\mu\right) = Y_{2H} - \left[i^*\left(\mu\right)h\left(V_H^{LCS}\left(y_H\right)\right) + \int_{i^*(\mu)}^1 h\left(V_H^o\left(s_i\left(\mu\right)\right)\right)di\right].$$

If $\mu = 1$, the result is trivial as that corresponds to full information disclosure. If $\mu < 1$,

$$-u'\left(C_1\right)\int_{i^*(\mu)}^1 \partial_\mu h\left(V_H^o\left(s_i\left(\mu\right)\right)\right)di + \int_{i^*(\mu)}^1 \partial_\mu V_H^o\left(s_i\left(\mu\right)\right)di \leqslant 0.$$

Since

$$u'\left(C_1\right)h'\left(V_H^o\left(s_1\left(\mu\right)\right)\right)\int_{i^*(\mu)}^1 \partial_\mu V_H^o\left(s_i\left(\mu\right)\right)di \geqslant u'\left(C_1\right)\int_{i^*(\mu)}^1 \partial_\mu h\left(V_H^o\left(s_i\left(\mu\right)\right)\right)di$$

by convexity of h, then

$$u'(C_1)\, h'(V_H^o(s_1(\mu))) \geqslant 1.$$

Let μ^* be the optimal signal structure in the economy without discrimination. Suppose by contradiction that $\mu^* > \mu$. Then

$$u'(C_1(\mu)) \geqslant \frac{1}{h'(V_H^o(s_1(\mu)))} > \frac{1}{h'(V_H^o(s_1(\mu^*)))} = u'(C_1^*(\mu^*))$$

where the second inequality follows by convexity of h and the last equality from interior optimality since $\mu^* > 0$. It follows that $C_1(\mu) < C_1^*(\mu^*)$ or, equivalently, $\Pi(\mu) < \Pi^*(\mu^*)$. Since $V_H^o(s_i(m)) \leqslant V_H^o(s_1(m))\ \forall i$, then $\Pi(\tilde{\mu}) \geqslant \Pi^*(\tilde{\mu})\ \forall \tilde{\mu}$, which implies $\Pi^*(\mu) < \Pi^*(\mu^*)$. Since Π^* is a decreasing function, the latter inequality constitutes a contradiction. Finally, if $V_H^o(s(m)) > V_H^{LCS}$, then $\Pi(\mu) > \Pi^*(\mu)$ which implies that $\mu > \mu^*$. Q.E.D.

E Effort

In this section, we extend the benchmark model to incorporate unobservable effort in the first period of the contract. In order to highlight the effect of this extension on optimal information design, we intentionally keep the environment as close as possible to our benchmark.

E.1 Environment

Consider a model in which first period output y_1 is constant and predetermined. At the end of the first period, the consumer exerts effort e at cost $v(e)$. At the beginning of the second period, the incumbent and the consumer jointly observe the outcome of the effort in the form of productivity $z \in \{z_H, z_L\}$ where $z \sim f(z|e)$. Effort is an investment-like good that does not affect first period output, but contributes to the formation of general productivity. Productivity in turn affects the distribution of second subperiod output, $y_2 \in \{y_L, y_H\}$ with $y_2 \sim p(y_2|z)$. We assume that while effort is privately known only by the consumer, productivity is observed by both the consumer and the incumbent, but not by outsiders. Outsiders only observe a signal m about the value of productivity, $m \sim \mu(z)$.

An allocation is a contract offered by the incumbent, $\{C_1, e, C_2(z, m)\}$, and a menu of contracts offered by the outsiders.

Additional assumptions and definitions We define $E(y_2|z) = \sum_s p(y_s|z)\, y_s$, with $E(y_2|z_H) > E(y_2|z_L)$, and $E(y_2|e) = \sum_z f(z|e)\, E(y_2|z)$. To make the environment comparable to the

economy in the text, we assume that y_1 is equal to Y_1 in the benchmark model. We also assume that $E(y_2|z_s) = Y_{2s}$ and $E(y_2|e) = Y_2$ which requires us to specify a given level of effort. Hence we assume that effort can take on two values, e and 0, and in equilibrium we guess (and verify) that is optimal to exert effort e. Furthermore, we assume that $f(z_H|e) = \pi_1(y_H)$ and $p(y_2|z_s) = \pi_2(y_2|y_s)$ for $s = H, L$.

In this economy, productivity replaces first period income as the source of information about future income that both the consumer and the incumbent have access to and that the designer conveys a signal about. The key difference with our benchmark model is that the determination of productivity z is influenced by an action the consumer performs and the incumbent cannot observe.[27] Thus, the two economies are equivalent except for the existence of an incentive compatibility constraint that guarantees that the consumer exerts effort e.

Equilibrium and optimal public disclosure We solve for the equilibrium under a given signal structure. Most of the results follow directly from what we showed in the text, hence we focus on the new features that originate from the introduction of effort. We discuss the economically interesting case in which it is efficient to induce strictly positive effort. Due to adverse selection, a consumer with a low realization of productivity always consumes her expected output, $E(y_2|z_L)$, which is also the only contract the outsiders are willing to offer. However, in order to induce the outsiders not to offer a contract that would attract borrowers with z_H, the incumbent offers them

$$V(z_H, m) = V_H^o(s(z_H|m; e))$$

where $s(z_H|m; e)$ is the share of agents with productivity z_H with signal m given the equilibrium effort level e,

$$s(z_H|m; e) = \frac{\mu(m|z_H) f(z_H|e)}{\sum_z \mu(m|z) f(z|e)}.$$

The key difference with the previous model is the endogeneity of productivity. For the borrower to exert the effort level e, it has to be that

$$(f(z_H|e) - f(z_H|0))[V(z_H, m) - u(Y_{2L})] \geqslant v(e) - v(0)$$

[27] In a labor market application, we can think of human capital as not fully revealing the amount of effort the worker exerts. The distribution f might be induced by either unobserved worker type or by pure luck. What matters for our results is that the source of uncertainty underlying f is uncorrelated with that behind p.

or

$$[V(z_H, m) - u(Y_{2L})] \geqslant \frac{v(e) - v(0)}{(f(z_H|e) - f(z_H|0))}. \tag{34}$$

If this incentive compatibility constraint is satisfied, the equilibrium consumption allocation is identical to the one in the benchmark model, otherwise borrowers exert no effort.

Notice that the information design has no effect on the value received by borrowers with low productivity. The only way to induce effort is to provide additional information about borrowers with high productivity, hence increasing their equilibrium value. We summarize this result in the following proposition:

Proposition 6. *Let μ^* be the optimal signal structure in the economy without hidden effort. i) If (34) holds, then μ^* is optimal in the economy with effort. ii) If (34) does not hold when evaluated at μ^*, then the optimal signal structure is: $M = \{g, b\}$ (good or bad) and $\mu(g|z_H) = 1$ and $\mu(b|z_L) \in (0, 1)$ such that*

$$V_H^o\left(\frac{f(z_H|e)}{f(z_H|e) + (1 - \mu(b|z_L))f(z_L|e)}\right) = \frac{v(e) - v(0)}{(f(z_H|e) - f(z_H|0))} + u(Y_{2L})$$

Moreover, $\mu(b|z_L) > \mu^*(b|z_L)$. That is, the optimal signal with hidden effort is more informative than the one without hidden effort.

Proof. Part $i)$ is straightforward since μ^* satisfies the IC constraint and it is optimal in its absence. Part $ii)$ follows from concavity of (10), the binding incentive constraint (34) and $V_H^o(s)$ being increasing in s. Q.E.D.

F Consumer saving

In the benchmark economy, the insurance contract is the only vehicle available to the consumer to transfer resources between period 1 and period 2. In this section, we allow the consumer to save at a gross rate equal to one, the rate at which they value resources across periods. We show that, under two-sided lack of commitment, the optimal public disclosure policy then reveals no information for every value of $\pi_1(y_H)$. Disclosure raises the period 2 transfers received by high-income consumers at the expense of period 1 transfers. When consumers can save, they can transfer resources from period 1 to both states in period 2, including the low-income state where marginal utility is highest; thus, it is optimal to maximize the incumbent's ex-post profits after $y_1 = y_H$, rebate them in period 1, and let consumers save to self-insure against a bad income realization.

Formally, in the first subperiod of period 1, the consumer can privately store $a \geqslant 0$ units of consumption, which pay a units of consumption in the first subperiod of period

2, regardless of the realization of y_1.[28] Since savings are chosen before income is realized, all consumers enter period 2 with the same savings a.[29]

Savings do not affect the equilibrium contract in period 2. They simply increase the consumer's expected consumption. Consider a consumer who enters period 2 with savings a. In the Miyazaki--Wilson program (2), the nonnegative profit condition involves only the resources of the firm, while the incentive compatibility and participation constraints compare the consumption that the same consumer obtains under different contracts, so the savings a add to both sides and cancel. Hence, $C_H^o(s)$ is independent of a, and the outside option value–including saving–is $u(C_H^o(s)+a)$ for the high type and $u(Y_{2L}+a)$ for the low type. The same argument applies to the retention offer of the incumbent and to the withdrawal stage: every comparison between contracts is unaffected by a common shift in consumption.[30] The characterization in Lemmas 4 and 5 therefore continues to hold, with the consumer now consuming C_1-a in period 1 and $C_H^o(s(m))+a$ in period 2 if $y_1=y_H$ and $Y_{2L}+a$ if $y_1=y_L$. The consumer's optimal saving a solves

$$\max_{a\geqslant 0} u(C_1-a)+\pi_1(y_H)\sum_m \mu(m|y_H)\,u(C_H^o(s(m))+a)+\pi_1(y_L)\,u(Y_{2L}+a). \tag{35}$$

Lemma 9. *Given a public disclosure policy (M,μ), the equilibrium consumption offered by firms is as in Lemma 5, and the consumer's optimal saving a solve (35).*

We now turn to the optimal disclosure policy.

Proposition 7. *Suppose that the consumer can save at a gross rate equal to one and that neither the firm nor the consumer can commit. Then, for every value of $\pi_1(y_H)$ the optimal public disclosure policy reveals no information.*

Proof. It remains optimal to assign the same signal to all high-income consumers: the argument in the proof of Proposition 2 relies only on the convexity of $C_H^o(s)$ and the concavity of u, and savings shift all period-2 consumption levels by the same constant. As in Section 5.1, the designer then effectively chooses the period-2 consumption of high-

[28] Storage within the period is never used: consumers value second-subperiod consumption at $\beta<1$, so moving consumption from the first to the second subperiod at rate one is strictly unprofitable. We also do not allow the consumer to store consumption received in the second subperiod of period 1.

[29] This implies that whether savings are hidden or not is not important. Firms in period 2 cannot make inference on the agent's type by observing their savings.

[30] This exact irrelevance result is driven by the linear aggregation of consumption within a period. In the economy studied in Dovis and Martellini (2024), savings will affect the incentive constraint in the Miyazaki--Wilson program. Yet, the conclusion of Proposition 7 below, that no information disclosure is optimal for every $\pi_1(y_H)$ when consumers can save, holds in that model as well.

income consumers, $C_{2H} = C_H^o(s(g))$, and the optimal disclosure policy solves

$$\max_{C_{2H} \in [C_H^o(\pi_1(y_H)), Y_{2H}]} W(C_{2H}) \tag{36}$$

where

$$W(C_{2H}) \equiv \max_{a \geqslant 0} u(C_1(C_{2H}) - a) + \pi_1(y_H)\, u(C_{2H} + a) + \pi_1(y_L)\, u(Y_{2L} + a) \tag{37}$$

and $C_1(C_{2H}) = Y_1 + \pi_1(y_H)(Y_{2H} - C_{2H})$, with the signal structure recovered from C_{2H} as in (17).

The objective in (37) is strictly concave in a, so the optimal savings a^* are unique and satisfy the necessary and sufficient first-order condition

$$u'(C_1(C_{2H}) - a^*) \geqslant \pi_1(y_H)\, u'(C_{2H} + a^*) + \pi_1(y_L)\, u'(Y_{2L} + a^*), \tag{38}$$

with equality if $a^* > 0$.[31] Moreover, for every C_{2H} in the admissible range,

$$C_{2H} \geqslant C_H^o(\pi_1(y_H)) \geqslant C_H^{LCS} = Y_{2L} + \beta\kappa(\pi_2(y_H|y_H) - \pi_2(y_H|y_L))(y_H - y_L) > Y_{2L}, \tag{39}$$

where the equality follows from the least-cost-separating allocation (4) by subtracting the binding incentive compatibility constraint of the low-income consumers, $Y_{2L} = c_{am} + \beta\pi_2(y_H|y_L)\, c_{pm}$, from the expected consumption of the high-income consumers, $C_H^{LCS} = c_{am} + \beta\pi_2(y_H|y_H)\, c_{pm}$, and using $c_{pm} = \kappa(y_H - y_L)$. Combining (38) with $u'(Y_{2L} + a^*) > u'(C_{2H} + a^*)$ yields

$$u'(C_1(C_{2H}) - a^*) > u'(C_{2H} + a^*). \tag{40}$$

Thus, consumption is strictly back-loaded after a high income realization, under every disclosure policy.

By the envelope theorem, W is differentiable with

$$W'(C_{2H}) = \pi_1(y_H)\,[u'(C_{2H} + a^*) - u'(C_1(C_{2H}) - a^*)] < 0,$$

where the inequality follows from (40). Thus, W is strictly decreasing and the unique optimum is $C_{2H} = C_H^o(\pi_1(y_H))$, attained by assigning the same signal to all consumers, $s(g) = \pi_1(y_H)$. Q.E.D.

[31] Saving the entire period-1 consumption is never optimal: at $a = C_1(C_{2H})$, consumption in period 1 is zero and strictly below consumption in period 2 in both states, so a marginal reduction in savings raises utility.

Remark. Savings play no role under one-sided commitment. Any allocation the consumer attains by saving can be replicated by the contract directly, so the optimal contract of Section 4 is unchanged. Moreover, the consumer has no incentive to save at the equilibrium allocation: by (9), $C_1 = C_2(y_L) \leqslant C_2(y_H)$, so

$$u'(C_1) \geqslant \pi_1(y_H)\, u'(C_2(y_H)) + \pi_1(y_L)\, u'(C_2(y_L))$$

and the first-order condition for savings holds at $a = 0$. With commitment, the firm back-loads consumption on the consumer's behalf, and Proposition 1 is unchanged.