Generous Long-Term Contracts

By

Sylvain Chassang, Princeton University

Griswold Center for Economic Policy Studies
Working Paper No. 354, May 2026

Generous Long-Term Contracts

Sylvain Chassang*

Princeton University

May 10, 2026

Abstract

I argue that in long-term consumer–producer relationships, menus of contracts can be advantageously replaced by a single generous contract such that, at any point in time, the sum of a consumer's transfers is equal to the sum of transfers they would have made under the contract that is the best for them in hindsight. Such generous long-term contracts can increase skeptical consumers' demand for complex and higher-powered contracts while approximately implementing the same outcomes as the underlying menu evaluated by a rational decision maker. Applications include voluntary load shedding in retail electricity markets and cost sharing in health insurance.

KEYWORDS: menus, load shedding, cost sharing, demand response, generous contracts.

1 Introduction

This paper argues that in long-term consumer-producer relationships, menus of contracts can often be profitably replaced by a single generous contract that guarantees consumers that at every time horizon, the sum of payments they have made is equal to the sum of payments they would have made under the contract that is the best for them in hindsight. Such

*Contact: chassang@princeton.edu. I'm grateful to Laurent Bozzi, Rebecca Dizon-Ross, Claire Del Gatto, Martin Vaeth, Isabelle Vautier, and Charles Weymuller for early motivating discussions. The paper benefited from feedback from Liran Einav, Ben Handel, Alessandro Lizzeri, Robert Metcalfe, Juuso Toikka, and Owen Zidar.

contracts, referred to as *generous long-term contracts* are obviously beneficial to consumers, potentially addressing a boundedly rational reluctance to adopt more complex and more high powered contracts. The surprising finding is that in many environments, they approximately implement the outcomes achieved under the baseline menu evaluated by rational decision-makers.

The case of load shedding in retail electricity markets illustrates the potential benefits of generous long-term contracting. Because electric utilities are regulated monopolies, they are often mandated to offer an affordable fixed price contract. However, such contracts give consumers no incentives to modulate their demand as a function of excess supply or demand. This leads electric utilities, such as French electric utility EDF, to introduce variable price contracts that charge consumers significantly higher prices at times when the grid is under strain.[1] However, demand for these variable price contracts is often low, even if they are attractively priced. In the case of a variable price contract offered by EDF, overall adoption is around 3% of households even though the overwhelming majority of consumers would benefit *without changing their consumption choices.* The promise of generous long-term contracting is to enhance the adoption of variable price contracts, without sacrificing the efficiency gains they generate, potentially allowing utilities to price them less attractively while increasing take up. A similar insight applies to the adoption of cost sharing contracts in health insurance.

The paper considers a dynamic principal-agent problem without discounting in which every period t, a consumer (the agent) observes a public state k_t, a private preference parameter α_t, and makes a consumption choice q_t. The consumer makes payments to a producer (the principal) according to one of two contracts of the form $\tau(q_t, k_t)$. A baseline contract τ^0 serves as reference, while a new contract τ^1 (e.g. a variable price contract for retail electricity markets) is being introduced by the producer. If contract τ^1 does not dominate reference

[1] EDF was the world's largest electricity generator in 2009. Its contractual offering, as well as a synthetic panel of individual demand data representative of French consumption patterns, will be used to calibrate performance metrics for generous long-term contracts.

contract τ^0 for the consumer, then the consumer evaluates the resulting menu $\mathcal{M} \equiv \{\tau^0, \tau^1\}$ using an adversarial (i.e. ambiguity averse) stance. When τ^1 dominates τ^0, the consumer evaluates menu $\mathcal{M} \equiv \{\tau^0, \tau^1\}$ using a neutral stance, i.e. under the prior shared ex ante with the producer. This creates a very sharp trade-off between adoption (of contract τ^1) and efficiency gains given adoption. The paper studies the extent to which generous long-term contracts can improve this trade-off.

It is immediate that generous long-term contracts would be adopted since they mechanically dominate reference contract τ^0. The surprising result is that when preferences α_t and state k_t are i.i.d. conditional on a consumer's type, then generous long-term contracts approximately implement the same allocation as the original menu evaluated by a rational expected utility maximizing decision-maker. The result extends as is to non-stationary environments exhibiting fast-learning, i.e. environments in which the consumer rapidly learns which of the underlying contracts τ^0 and τ^1 is likely to serve them best.

In non-stationary environment with slow learning, i.e. environments in which the consumer remains uncertain over which contract will serve them best for a considerable amount of time, generous long-term contracts need not implement the same allocation as the underlying menu evaluated by a rational consumer. However, it is possible to provide a partial extension for a class of alternative contracts τ^1 referred to as Pareto alignments. A contract τ^1 is a Pareto alignment of τ^0 if it is a weighted average of τ^0 and the producer's change in profit $\Delta\pi(q_t, k_t) \equiv \pi(q_t, k_t) - \overline{\pi}^0(k_t)$ where $\overline{\pi}^0(k_t)$ is an upper bound to the producer' expected counterfactual profit under τ^0. When τ^1 is a Pareto alignment of τ^0, then menu $\mathcal{M} = \{\tau^0, \tau^1\}$ necessarily yields a Pareto improvement over the single reference contract $\{\tau^0\}$. In addition, Pareto alignments satisfy a tight lower bound on the producer's profit increases given an increase in consumer welfare. When τ^1 is a Pareto alignment of τ^0, generous long-term contracts based on menu $\mathcal{M}$ yield an approximate Pareto improvement over contract τ^0, and approximately achieve the same profit guarantee for the producer, up to a penalty on per-period profits of order $1/\sqrt{N}$ where N is the number of periods in the contracting relationship.

Finally, the paper illustrates the empirical relevance of generous long-term contracts by calibrating the penalties associated with generous Pareto alignments in the context of load-shedding in the French retail electricity market, and cost-sharing in health insurance. The penalties are plausibly manageable given potential efficiency gains in both contexts. This calibration exercise yields two design insights. First, there is a trade-off between potential efficiency gains and profit penalties associated with generous treatment. Setting a high counterfactual profit estimate $\overline{\pi}^0(k_t)$ reduces penalties associated with generosity, but can eliminate potential efficiency gains associated with the underlying menu $\mathcal{M}$. Setting a counterfactual profit estimate closer to unconditional expected counterfactual profits can increase the potential efficiency gains under menu $\mathcal{M}$ but can result in large generosity penalties when consumers exhibit persistent differences in consumption. The second insight is that the intensity of this trade-off is mediated by the ability to capture persistent differences in the profitability of individual consumption choices in counterfactual profitability estimates. Provided that persistent individual differences can be accounted for, generous contracts can achieve the majority of welfare gains associated with Pareto alignments.

The paper contributes to the recent theoretical literature that seeks to develop methods to make contract theory better suited to practical implementation (Rogerson, 2003, Chassang, 2013, Carroll, 2015, Madarász and Prat, 2017, Carroll, 2017). Like Chassang (2013) and Chassang and Kapon (2022) it exploits error averaging properties of long-term contracts to relax constraints that bind in the case of short term contracts. In Chassang (2013) and Chassang and Kapon (2022) the constraints are limited liability constraints. The current paper imposes that new contracts should dominate reference contracts for the consumer.

The decision to take as given an initial menu $\mathcal{M}$ as a starting point to contract design reflects the *minimalist design* view advocated in Sönmez (2023) and Greenberg et al. (2024). Under the traditional mechanism design approach, the economist understands the underlying environment, the objectives, and has free-reign to design choice problems and incentive contracts. In high-stakes policy settings it is rarely possible to design incentives from scratch. Stakeholders and decision-makers have a vague understanding of the underlying environment that is embodied in imperfect solutions and coarse design principles. In order to actually effect change a minimalist designer does not seek to solve the problem from scratch, but rather seeks to formalize existing design insights, and makes marginal improvements consistent with those insights. The current paper takes as given the underlying menu whose performance generous long-term contracts seeks to replicate in a manner that is more appealing to skeptical consumers.

Finally, the paper seeks to contribute to the applied literature studying the adoption of better aligned cost-sharing contracts in sectors such as retail electricity, and health insurance markets (Allcott, 2011, Wolak, 2011, Jessoe and Rapson, 2014, Brot-Goldberg et al., 2017, Fowlie et al., 2021, Ho and Lee, 2023, Dizon-Ross and Zucker, 2025). In particular, Fowlie et al. (2021) studies the impact of setting variable price contracts as the default choice to increase the adoption of variable price electricity contracts. It shows that defaults can be extremely helpful, even though additional consumers have a lower elasticity of demand to price. The current paper complements this view by showing that in fact, it may not be necessary to offer consumers a menu in the first place.

Ito et al. (2023), Ida et al. (2025) study the design of menus of retail electricity contracts, where the main question is how much to subsidize the adoption of variable price contracts. They show that menus are helpful and that the benefits of subsidies are decreasing: consumers that value variable pricing more are also the one whose behavior is most affected. In addition, there is considerable value in indexing menus on observable consumer characteristics. Ida et al. (2024) studies two-period menu design, emphasizing the role of dynamic treatment effects and information acquisition to better target in future periods. The current paper highlights the value of generous long-term contracting as a way to reduce the magnitude of subsidies needed for consumers to adopt variable price contracts.

The paper is structured as follows. Section 2 provides a motivating example in the context of retail electricity markets, highlighting the potential upside of menus, and adoption challenges associated with variable price contracts. Section 3 introduces the framework. Section 4 studies generous long-term contracting in conditionally i.i.d. environments, while Section 5 turns to the case of non-stationary environments. Section 6 concludes with a calibration of profit penalties in the context of retail electricity markets, and cost-sharing in healthcare insurance.

2 Motivating Application: Load Shedding

Load shedding, sometimes referred to as demand response, corresponds to strategies electric grid operators use to adjust demand to match electricity supply fluctuations (Borenstein and Holland, 2005, Joskow and Wolfram, 2012). It is an important tool for grid management, especially given the recent growth of less pilotable renewable energy sources (Wolak, 2019, Reguant and Wagner, 2025).

One strategy consists in varying the hourly cost of electricity as a function of supply. There are two limits to the effectiveness of this strategy:

(i) because electricity providers are regulated utilities they are often required to offer a fixed-price contract at a regulated price.

(ii) because short-run demand is inelastic, large price variation is needed in order to induce a change in behavior (Reiss and White, 2005).[2]

As a result of (i), electric utilities end up offering a menu consisting of both fixed price and variable price contracts. Because of (ii) the price variation needed to change consumption behavior is large. This makes variable price contracts seem particularly risky, so that few

[2]This being said, Jessoe and Rapson (2014) show that providing consumers with easily accessible information about their electricity consumption more than doubles the price elasticity of retail electricity consumption.

consumers opt for them. For instance, French utility EDF offers a varying price contract branded as *Tempo* under which the utility can label 22 days a year as *Red Days*, and increase electricity prices by a factor of three (Aubin et al., 1995, Cabot and Villavicencio, 2024, RTE, 2025). The utility specifies a day ahead whether the next day is red or not, with the constraint that there can be at most five consecutive red days.[3] Although the variable price contract is attractively priced (electric optimization startup Lite (2023) reports that in its sample of clients, 100% of households would reduce their yearly consumption costs by approximately 20% without changing their consumption if they switched to Tempo contracts) national demand for the contract is modest: roughly 3% of households. As Figure 1 highlights, red days correspond to days when the hourly spot price tends to be above the base price charged by the utility.

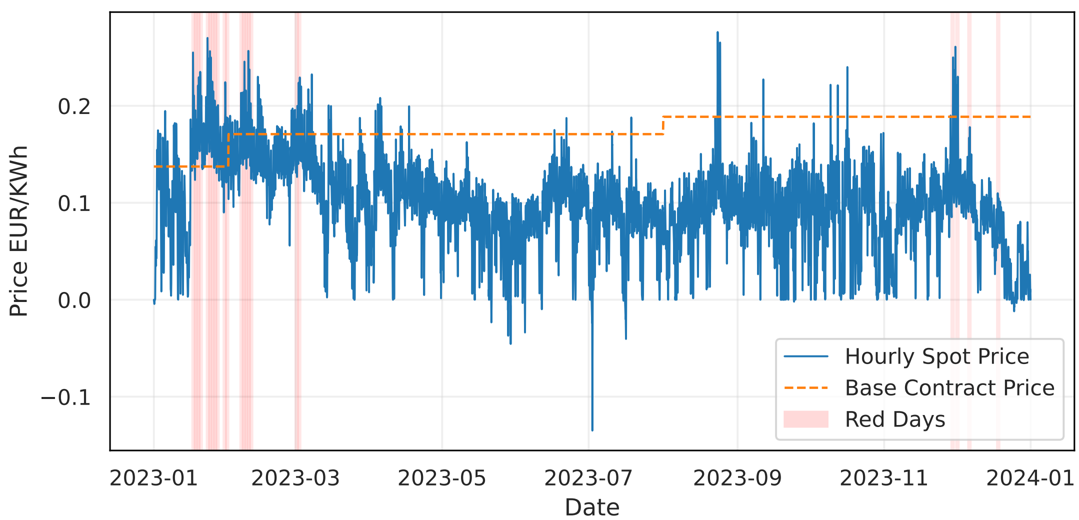

Figure 1: Spot prices, base contract prices, and red days (EDF Tempo, 2023)

The observation that menus of contract can enhance the payoffs of the principal has been well understood at least since Laffont and Tirole (1986). The rest of this section offers a simple one-period model to illustrate why such menus can be useful in the context of

[3]On average, variable price clients reduce their demand by 23% on high price days relative to low price days (EDF, 2025).

retail electricity markets, and why demand for variable price contracts may be low, even if they are attractively priced. The next sections study ways to relax the trade-off between implementation and adoption in long-term contracting environments.

2.1 A toy model

A consumer chooses a quantity q to consume. The cost c of producing q depends on state $k \in \{\underline{k}, \overline{k}\}$, with $\overline{k} > \underline{k}$, and takes the form $c(q, k) = kq$. The state k is equal to $\overline{k}$ with probability $a \in (0, 1)$.

Let us denote by $\tau \geq 0$ the transfer made by the consumer to the producer. The consumer's von Neumann-Morgenstern preferences over pairs (q, τ) take the form

$$v(q, \tau) \equiv \alpha \times (q\mathbf{1}_{q \leq q_0} + q_0\mathbf{1}_{q > q_0}) - \tau.$$

The consumer has an unconstrained ideal consumption equal to q_0, and a separable linear cost of transfers.[4] For simplicity, parameter q_0 is fixed, but preference parameter $\alpha \in \{\alpha_L, \alpha_H\}$, with $\alpha_L < \alpha_H$ can vary. Parameter α is equal to α_H with the same probability a that marginal costs of production k are equal to $\overline{k}$.

Consumers can be one of two types $\theta \in \{\theta_{corr}, \theta_{indep}\}$ depending on whether the intensity of their need for electricity is correlated with production costs or not. A consumer of type θ_{corr} is such that $\alpha = \alpha_H$ if and only if $k = \overline{k}$, i.e. the consumer values consumption when it is expensive to provide. In contrast, a consumer of type θ_{indep} is such that

$$\mathsf{prob}(\alpha = \alpha_H | k = \overline{k}, \theta_{indep}) = a.$$

In practice, types θ_{corr} may correspond to consumers using electric heating, while θ_{indep} corresponds to consumers using wood, or gas furnaces.

[4]Consumption q may be interpreted as additional consumption above a minimum level $\underline{q}$ that the consumer cannot go below.

Menus. Contracts $\tau(q,k)$ index transfers on observable consumption q and a publicly observable production state k. The producer is required to offer a fixed priced contract

$$\tau^0(q,k) \equiv k_0 q,$$

but is allowed to also offer a variable price contract τ^1:

$$\tau^1(q,k) \equiv \begin{cases} \overline{p}q & \text{if} \quad k = \overline{k} \\ \underline{p}q & \quad \text{otherwise} \end{cases} \qquad \text{with} \quad k_0 \in (\underline{p}, \overline{p}).$$

Let $\mathcal{M} = \{\tau^0, \tau^1\}$ denote the menu of contracts available to the consumer. Assume that the consumer picks a contract after they learn their type, but before the state k is revealed.

Timing. The timing of choices is as follows. The principal offers a menu $\mathcal{M}$. Given type θ, the consumer chooses a contract τ from the menu. A state of the world k and a preference α are then realized, and the consumer makes a consumption decision q leading to transfers $\tau(q,k)$.

Note that menu $\mathcal{M}$ would be ineffective at screening different types of consumers if menu choice took place after state k was realized: regardless of consumption level q, contract τ^0 is cheaper than τ^1 if $k = \overline{k}$, and more expensive than τ^1 if $k = \underline{k}$.

Consumption choices. Consider a linear contract $\tau(q) = pq$ with $p \in (0, \alpha)$. Optimal consumption q^* under that contract takes the form

$$q^* = \begin{cases} q_0 & \text{if} \quad \alpha > p \\ 0 & \quad \text{otherwise.} \end{cases}$$

This is associated with equilibrium utility

$$u^*(p|\alpha) = [\alpha - p]^+ q_0,$$

where $[x]^+ \equiv \max(x, 0)$ for all $x \in \mathbb{R}$.

The following assumption simplifies the analysis, and is maintained for the rest of this section.

Assumption 1. *(i)* $a\overline{p} + (1-a)\underline{p} = a\overline{k} + (1-a)\underline{k} = k_0$.

(ii) $\alpha_L < \underline{k} < \underline{p} < k_0 < \alpha_H < \overline{p} < \overline{k}$,

Point (i) requires that k_0 be equal to the average cost of production, while point (ii) requires that variable prices are a mean preserving spread of fixed-price k_0. Point (ii) also implies that demand will be 0 under both contracts when needs are low ($\alpha = \alpha_L$), and under the variable price contract when costs are high ($k = \overline{k}$).

2.2 The value of menus

This section shows that introducing menu $\mathcal{M}$ is a Pareto improvement over offering the fixed price contract $\{\tau^0\}$ alone.[5]

Consumption and welfare under fixed price. Assumption 1 implies that under fixed price contract τ^0 both types of consumer only consume when they have a high need α_H for the product, in which case they demand q_0. This implies an average demand aq_0 under both types, and expected payoff $aq_0(\alpha_H - k_0) > 0$ for both consumer types. The producer makes a profit $aq_0(k_0 - \overline{k}) < 0$ on the correlated-need types θ_{corr}, and a profit $aq_0(k_0 - a\overline{k} - (1-a)\underline{k}) = 0$ on the independent-need types θ_{indep}.

Consumption and welfare under variable price. Under the variable price contract the correlated-need consumer would always demand $q = 0$, while the independent-need consumer would demand $q = q_0$ only when $\alpha = \alpha_H$ and $k = \underline{k}$.

[5]If dollar-valued surplus maximization is the objective, then the first-best can be achieved without menus: offer electricity at the spot price. The constraint that contract τ^0 be offered may captures an implicit redistributive objectives along the lines of Dworczak et al. (2021). The objective of finding a Pareto improvement on this contract without questioning its appropriateness reflects the minimalist design approach of Sönmez (2023).

Note that since variable prices are a mean preserving spread of fixed price k_0, the independent-need consumer can guarantee itself the same payoff as under the fixed price contract by demanding q_0 independently of k. It follows from this that the correlated-need type θ_{corr} is strictly better under the fixed-price contract τ^0, while the independent-need type θ_{indep} is strictly better off under the variable price contract.

Finally, observe that the producer benefits from having independent-need consumers curtail their consumption, since profits from such consumers become positive: $a(1-a)q_0(\underline{p}-\underline{k}) > 0$. Altogether this implies the following.

Lemma 1 (Menus are useful)**.** *Under menu $\mathcal{M} = \{\tau^0, \tau^1\}$, the correlated-need type continues to use contract τ^0, while the independent need type signs up for contract τ^1. The welfare of the correlated-need consumer is unchanged, while the welfare of the independent-need user and the profits of the producer both increase.*

2.3 Why is demand for variable price contracts low in practice?

Adversarial inference. A subjectively plausible reason why the variable price contract is unattractive is that if the principal possesses private information about state k and preferences α, then the offer of a variable price contract τ^1 is adversely selected. This is the motivating observation of the literature on informed principals (Myerson, 1983, Maskin and Tirole, 1990, 1992).

The paper builds on this intuition using the following stark model. Let $u^*(\tau|\alpha, k)$ denote the equilibrium welfare of a consumer with preferences α, in state k under contract τ. The producer can be of two types:

(i) *Neutral*, in which case the producer has the same beliefs as the consumer regarding state k, is uninformed about consumer type α, and simply seeks to maximize profits.

(ii) *Adversarial*, in which case the producer knows the state k, knows the consumer's type α, and seeks to minimize the consumer's welfare $u^*(\tau|\alpha, k)$.[6]

Definition 1. *A contract τ^1 weakly dominates τ^0 (for the consumer) if for all (α, k), $u^*(\tau^1|\alpha, k) \geq u^*(\tau^0|\alpha, k)$.*

The consumer's beliefs following the introduction of contract τ^1 are such that:

(i) If τ^1 is weakly dominant, then the consumer believes the producer is neutral.

(ii) If τ^1 is not weakly dominant, then the consumer believes the producer is adversarial. The consumer updates their belief over their own preferences α and state k reflecting the knowledge that introducing contract τ^1 necessarily benefits the adversarial producer.[7]

Under this framework, because $\overline{p} > k_0 > \underline{p}$, demand for the variable price contract will be null: the producer is offering this contract because it expects prices to be high.

Generous short-term contracts. Even under adversarial inference, not every contract τ^1 is rejected. Weakly dominant contracts will be evaluated under the usual uninformed principal framework, i.e. the neutral stance. The challenge is that there is no more screening: the consumer may as well pick the dominant contract.

Starting from any desired menu $\mathcal{M} = \{\tau^0, \tau^1\}$, such that neither contract weakly dominates the other, a generic way to replace menu $\mathcal{M}$ is to consider the generous contract

$$\widehat{\tau}(q, k) \equiv \min\{\tau^0(q, k), \tau^1(q, k)\}.$$

[6]In this model, the adversarial producer is not only informed, but also out to get the consumer. This reflects two psychologically realistic features of consumer choice. First, consumers often take a zero-sum view of contractual relations with large corporations. Second decision-makers tend to evaluate evaluate complex high-powered contracts using ambiguity averse preferences (de Clippel et al., 2024).

[7]A natural extension following the formulation of Huber (1964) would have the consumer evaluate menus using a weighted average of expected and adversarial beliefs.

This contract weakly dominates τ^0, so that it would be picked in the event that menu $\widehat{M} \equiv \{\tau^0, \widehat{\tau}\}$ is offered. However, it is obvious that this generous contract does not implement the same allocation as the original menu $\mathcal{M} = \{\tau^0, \tau^1\}$ evaluated under the neutral stance.

Indeed, let us return to the context of linear contracts τ^0 and τ^1 satisfying Assumption 1, with τ^1 a mean-preserving spread of τ^0. The consumption choices and transfers respectively associated with menus $\mathcal{M} = \{\tau^0, \tau^1\}$ and $\widehat{\mathcal{M}} = \{\tau^0, \widehat{\tau}\}$ are as follows:

- For correlated-need consumers, consumption and transfers are unchanged whether consumers choose from menu $\mathcal{M}$ or $\widehat{\mathcal{M}}$.

- For independent-need consumers, consumption and transfers are the same for menus $\mathcal{M}$ and $\widehat{\mathcal{M}}$ in the low cost state $\underline{k}$, but they are different in the high cost state $\overline{k}$. Under menu $\mathcal{M}$ independent-need consumers do not consume in the high cost state, while they consume whenever $\alpha = \alpha_H$ under menu $\widehat{\mathcal{M}}$.

Altogether, the producer is worse off under $\widehat{\mathcal{M}}$ than under τ^0 since menu $\widehat{\mathcal{M}}$ leads to the same consumption patterns but lowers transfers. In other terms there is a real trade-off between increasing adoption (by offering a generous contract), and producer profits under the adopted contract.

The rest of the paper argues that this trade-off vanishes as the duration of contracting relationships increases.

3 General Framework

Uncertainty, consumption, and preferences. Consider a dynamic consumption problem, with finite horizon N and discrete time $t \in \{1, \ldots, N\}$. Ahead of any choice, the consumer is informed of their type $\theta \in \Theta$, with Θ finite. Types are fully persistent. In each period t:

1. a publicly observable, and contractible state $k_t \in K$ is drawn; k is exogenous to the consumer's behavior but can include supply side factors (e.g. weather, spot prices, population level medical spending) as well a client specific characteristics (e.g. type of heating, housing, age, medical diagnostics);

2. a preference parameter $\alpha_t \in A$ is drawn and privately observed by the consumer;

3. the consumer makes a consumption choice $q_t \in Q$.

For simplicity, K, A and Q are assumed to be finite. The consumer's vNM preferences over streams of consumption $\mathbf{q} \equiv (q_t)_{t\in\{1,\ldots,N\}}$ and transfers $\tau \equiv (\tau_t)_{t\in\{1,\ldots,N\}}$ take the form

$$\frac{1}{N}\sum_{t=1}^{N} u(q_t, \alpha_t) - \tau_t.$$[8]

Time is not discounted, and preferences over consumption and transfers are separable. This framework is best suited to model behavior over moderate time-horizons corresponding to at most a few years. Section 6 proposes a calibration in the context of retail electricity using a one year horizon.

In each period t the firm experiences a production cost $c(q_t, k_t)$, and values streams of transfers and consumption according to

$$\frac{1}{N}\sum_{t=1}^{N} \tau_t - c(q_t, k_t).$$

Let μ denote the consumer and the producer's common prior over types and sequences of preferences and states $(\theta, \alpha_t, k_t)_{t\in\{1,\ldots,N\}}$.

[8] Appendix A extends the analysis to preferences of the form $\frac{1}{N}\sum_{t=1}^{N} u(q_t, \alpha_t) - \Gamma\left(\frac{1}{N}\sum_{t=1}^{N} \tau_t\right)$, with Γ increasing and convex. For instance, if the consumer faces a budget constraint, Γ may be linear over a range $[0, \overline{T}]$ and equal to $+\infty$ for expenses above $\overline{T}$.

Contracts. The principal is constrained to offer a given static contract $\tau^0 : (q_t, k_t) \mapsto \tau^0_t \in \mathbb{R}$, and is considering offering an alternative static contract $\tau^1 : (q_t, k_t) \mapsto \tau^1_t \in \mathbb{R}$. However, contract τ^1 does not weakly dominate τ^0 for the consumer so that demand is low.[9]

Definition 2 (generous long-term contracts)**.** *Let $h_t = (k_s, q_s)_{s\in\{1,\dots,t\}}$ denote the public history up to the end of period t. A long-term contract τ is a mapping $\tau : h_t \mapsto \tau(h_t) \in \mathbb{R}$ from public histories to contemporaneous transfers.*

A long-term contract τ is generous if it progressively weakly dominates benchmark contract τ^0: for all public histories h_t,

$$\sum_{s=1}^{t} \tau(h_s) \leq \sum_{s=1}^{t} \tau^0(q_s, k_s).$$

Evaluation stances. As in Section 2, we contrast two evaluation stances for menus, implicitly reflecting inference from menu expansions by a skeptical consumer.

Definition 3 (evaluation stances)**.** *Menu $\mathcal{M}$ is evaluated according to a* neutral stance*, if the consumer evaluates contract choice assuming the producer shares their prior μ.*

Menu $\mathcal{M}$ is evaluated according to an adversarial stance*, if the consumer assumes that the producer anticipates future states, and consumption choices, and seeks to maximize the consumer's sum of transfers.*

Assumption 2 (adversarial inference)**.** *Menus $\mathcal{M}$ that expand on τ^0 with generous contracts are evaluated under a neutral stance, while menus $\mathcal{M}$ that expand on τ^0 with non-generous contracts are evaluated under an adversarial stance.*

In addition, the paper considers two paradigms for contract choice:

- Under *repeated choice*, the consumer gets to pick a contract every period t, after observing type θ, and the history of preference and state realizations $(\alpha_s, k_s)_{s<t}$, before preference parameter α_t and state k_t are realized.

[9] The analysis extends essentially as is to finite menus with more than 2 options.

- Under *single choice*, the consumer gets to pick once and for all from menu $\mathcal{M}$ at the beginning of period $t = 1$ knowing only their type θ.

The single choice paradigm is practically the most relevant one since, in practice, contract changes are infrequent. Still the fact that performance guarantees can extend to the repeated choice framework is reassuring, since consumers can shop for contracts at a cost. In addition, the repeated choice case can be useful as a technical device to derive performance bounds for the single choice case.

The main question of interest can be reformulated as follows: does there exists a generous contract that approximately implements the outcomes the non-generous menu $\mathcal{M} \equiv \{\tau^0, \tau^1\}$ would achieve under a neutral stance?

Strategies. A private history h_t^α takes the form $h_t^\alpha \equiv (\alpha_s, k_s, q_s)_{s \leq t}$. Since the consumer is an expected utility maximizer under the neutral stance, it is without loss of generality to focus on pure dynamic strategies.

Under the single choice paradigm, a dynamic strategy for the consumer is a pair (τ, σ), with $\tau \in \{\tau^0, \tau^1\}$, and $\sigma \in \Sigma$ a consumption strategy that maps the last period history h_{t-1} to a current action plan $\sigma(h_{t-1}) \in Q^{A \times K}$. Σ denotes the set of dynamic consumption strategies.

Under the repeated choice paradigm, a dynamic strategy is a pair (ϕ, σ) where $\phi \in \Phi$ is a dynamic contract choice strategy that maps the last period history h_{t-1} to a current contract $\phi(h_{t-1}) \in \{\tau^0, \tau^1\}$, while $\sigma \in \Sigma$ is a dynamic consumption strategy as described above. To unify notation, let $\Phi = \mathcal{M}$ under the single choice paradigm.

4 Conditionally I.I.D. Environments

This section investigates generous long-term contracting in the class of environments where pairs of public states and private preferences (k_t, α_t) are jointly i.i.d. conditional on consumer type θ.

The generous long-term contract $\widehat{\tau}$ is defined as follows: for any public history h_t, let

$$\widehat{T}(h_t) \equiv \min_{\tau \in \{\tau^0, \tau^1\}} \sum_{s=1}^{t} \tau(q_s, k_s),$$
$$\text{and} \quad \widehat{\tau}(h_t) \equiv \widehat{T}(h_t) - \widehat{T}(h_{t-1}). \tag{1}$$

A useful observation is that flow payments under $\widehat{\tau}$ are bracketed by flow payments under either of the underlying contracts. This is practically important: this means that there is no need for large “catch-up" payments if it turns out that a consumer initially planning to be charged according to contract τ^1 ends up being charge according to contract τ^0.

Lemma 2 (bounded flow payments)**.** *For any public history h_t,*

$$\min_{\tau \in \{\tau^0, \tau^1\}} \tau(h_t) \leq \widehat{\tau}(h_t) \leq \max_{\tau \in \{\tau^0, \tau^1\}} \tau(h_t).$$

4.1 A heuristic

Section 2 showed that in short-term settings, the ex post lowest cost contract $\widehat{\tau}$ fails to approximately implement the outcomes of menu $\mathcal{M} = \{\tau^0, \tau^1\}$ under a neutral stance. The following heuristic argument suggests why this approach can plausibly succeed in long-term settings.

Consider the single-choice paradigm.[10] A consumer of type θ solves

$$\max_{\tau \in \mathcal{M}} \max_{\sigma \in \Sigma} \mathbb{E}_{\alpha, k|\theta} \left[\frac{1}{N} \sum_{t=1}^{N} u(q_t, \alpha_t) - \frac{1}{N} \sum_{t=1}^{N} \tau(q_t, k_t) \,\middle|\, \sigma \right] \tag{2}$$

The assumptions made so far imply that the optimal consumption choice is a function of

[10]The same argument holds in the repeated choice paradigm since optimal contract choice is constant in conditionally i.i.d. environments.

only current preferences α_t and state k_t. Let $\Sigma_Q^M = Q^{A\times K}$ denote pure Markovian strategies mapping preferences and states $A \times K$ to consumption choices Q.

The consumer's optimal choices given menu $\mathcal{M}$ under a neutral stance satisfy

$$\max_{\tau\in\mathcal{M}} \max_{\sigma\in\Sigma_Q^M} \mathbb{E}_{\alpha,k|\theta}\left[\frac{1}{N}\sum_{t=1}^{N} u(q_t,\alpha_t) - \frac{1}{N}\sum_{t=1}^{N}\tau(q_t,k_t)\,\Big|\,\sigma\right]$$
$$= \max_{\sigma\in\Sigma_Q^M}\left(\mathbb{E}_{\alpha,k|\theta}\left[\frac{1}{N}\sum_{t=1}^{N} u(q_t,\alpha_t)\,\Big|\,\sigma\right] - \min_{\tau\in\mathcal{M}} \mathbb{E}_{\alpha,k|\theta}\left[\frac{1}{N}\sum_{t=1}^{N}\tau(q_t,k_t)\,\Big|\,\sigma\right]\right).$$

Since Σ_Q^M is finite, the law of large numbers implies that uniformly over $\sigma \in \Sigma_M^Q$, the following approximate equality (suppressing the dependency of expectations on σ) holds:

$$\min_{\tau\in\mathcal{M}} \mathbb{E}_{\alpha,k|\theta}\left[\frac{1}{N}\sum_{t=1}^{N}\tau(q_t,k_t)\right] = \min_{\tau\in\mathcal{M}} \mathbb{E}_{\alpha,k|\theta}\left[\mathbb{E}_{\alpha,k|\theta}\left[\tau(q_1,k_1)\right]\right] = \mathbb{E}_{\alpha,k|\theta}\left[\min_{\tau\in\mathcal{M}} \mathbb{E}_{\alpha,k|\theta}\left[\tau(q_1,k_1)\right]\right]$$
$$\simeq \mathbb{E}_{\alpha,k|\theta}\left[\min_{\tau\in\mathcal{M}} \frac{1}{N}\sum_{t=1}^{N}\tau(q_t,k_t)\right].$$

Hence the consumer's optimal choices given menu $\mathcal{M}$ under a neutral stance approximately solve

$$\max_{\sigma\in\Sigma_Q^M} \mathbb{E}_{\alpha,k|\theta}\left[\frac{1}{N}\sum_{t=1}^{N} u(q_t,\alpha_t) - \frac{1}{N}\sum_{t=1}^{N}\widehat{\tau}(h_t)\,\Big|\,\sigma\right].$$

This provides the foundation for why generous contract $\widehat{\tau}$ may be a successful policy in long-term contracting environments. Sample averages become arbitrarily close to expectations which the consumer already optimizes over.[11] The next section provides a formal result. To do so, two technical points must be established:

- First, approximately correct incentives are sufficient to induce approximately correct choices. This can fail if global incentive compatibility constraints are binding.

- Second, the generous contract does not induce the consumer to use a non-stationary

[11]This point is loosely related to the observation in Azevedo and Budish (2019) that indexing mechanisms on endogenous outcomes need not meaningfully threaten strategy-proofness provided endogenous outcomes are aggregated over sufficiently large populations.

strategy under which the law of large numbers fails. This is not immediate since taking the minimum over contracts makes the consumer risk-loving with respect to total payments. This encourages the consumer to use consumption strategies under which average total payments remain a non-degenerate random variable.

4.2 Performance approximation

To simplify notation, for any flow contract τ, and sequence $(\alpha_t, k_t, q_t)_{t\in\{1,\ldots,N\}}$ of private preferences, public state, and consumption, let

$$U \equiv \frac{1}{N}\sum_{t=1}^{N} u(q_t, \alpha_t) \quad \text{and} \quad T_\tau \equiv \frac{1}{N}\sum_{t=1}^{N} \tau(h_t).$$

Any dynamic strategy (ϕ, σ) (where $\phi \in \Phi = \mathcal{M}$ in the single choice paradigm) is associated with an averaged-out mixed Markov strategy $\mathsf{e}[\phi, \sigma]$ defined by

$$\mathsf{e}[\phi, \sigma] \equiv \mathbb{E}\left[\frac{1}{N}\sum_{t=1}^{N}(\phi(h_{t-1}), \sigma(h_{t-1}))\right].$$

Note that in the repeated choice paradigm, $\mathsf{e}[\phi, \sigma]$ is a correlated Markov strategy: it belongs to $\Delta\left(\mathcal{M} \times Q^{A\times K}\right)$. In the single-choice paradigm, $\mathsf{e}[\phi, \sigma] \in \mathcal{M} \times \Delta\left(Q^{A\times K}\right)$. Similarly, let $\mathsf{e}[\sigma] \in \Delta\left(Q^{A\times K}\right)$ denote the expected time average of σ.

For simplicity, the remainder of this section maintains the following assumption.

Assumption 3. *For every type θ, mapping*

$$\nu \in \Delta\left(\mathcal{M} \times Q^{A\times K}\right) \mapsto \mathbb{E}_{q,\alpha|\nu,\theta}[U] - \mathbb{E}_{q,k|\nu,\theta}[T_\tau] \tag{3}$$

has a unique maximizer ν^*_θ.[12]

[12] When this assumption does not hold, approximately correct incentives do not imply approximately correct behavior. If this is the case, then both baseline contracts must be adjusted by placing a small amount of weight on producer profits along the lines of the Pareto alignments described in Section 5. As in

Note that since objective (3) is linear in ν, a unique maximizer is necessarily in pure strategies. This is not guaranteed in the extension considered in Appendix A, where the consumer is allowed to be risk-averse with respect to total transfers.

Under either the single or repeated choice paradigm, let $(W_{\mathcal{M}}(\theta), \Pi_{\mathcal{M}})_{\theta\in\Theta}$ and $(\widehat{W}(\theta), \widehat{\Pi}(\theta))_{\theta\in\Theta}$ denote expected consumer welfare and producer profits associated with each type θ given a neutral stance, respectively under menu $\mathcal{M} = \{\tau^0, \tau^1\}$ and generous long-term contract $\widehat{\tau}$.

Proposition 1 (performance approximation)**.** *Under either the single or repeated choice paradigm, for all $\theta \in \Theta$, the following hold:*

(i) $\widehat{W}(\theta) \geq W_{\mathcal{M}}(\theta)$*;*

(ii) for any $\eta > 0$, for N large enough, $\widehat{\Pi}(\theta) \geq \Pi_{\mathcal{M}}(\theta) - \eta$

Point (i) is immediate since consumers are offered a contract that weakly dominates both options in menu $\mathcal{M}$.

The proof of point (ii) is instructive. Take as given a type θ. The dependency of expectations on θ is suppressed for readability. Let $\widehat{\sigma}$ and $\sigma_{\mathcal{M}}$ respectively solve

$$\max_{\sigma\in\Sigma} \mathbb{E}_\sigma \left[U - \widehat{T}\right] \quad \text{and} \quad \max_{\sigma\in\Sigma} \left(\mathbb{E}_\sigma[U] - \min_{\phi\in\Phi} \mathbb{E}_\sigma \left[T_\phi\right]\right)$$

Note that $\widehat{\sigma}$ and $\sigma_{\mathcal{M}}$ implicitly depend on N. Recall that by convention, Φ degenerates to $\mathcal{M}$ under the single choice paradigm.

Let $||\cdot||$ denote the Euclidean distance on the set of mixed Markov strategies $\Delta(\mathcal{M} \times Q^{A\times K})$ viewed as a finite dimensional simplex. Recall that ν^*, defined by Assumption 3, belongs to the set of mixed Markov strategies.

Lemma 3 (strategy averages converge)**.** *Under either the repeated or single choice paradigm, as N grows large, the following hold*

Chassang (2013) and Madarász and Prat (2017), this ensures that if global IC constraints are tight in the original problem, they are broken in a direction that benefits the principal under approximate incentives.

(i) Markov averages $\mathsf{e}[\widehat{\sigma}]$ *and* $\mathsf{e}[\sigma_{\mathcal{M}}]$ *asymptotically solve*

$$\max_{\sigma \in \Delta(Q^{A \times K})} \mathbb{E}_{q,\alpha|\sigma,\theta}[U] - \min_{\tau \in \mathcal{M}} \mathbb{E}_{q,k|\sigma,\theta}[T_\tau].$$

(ii) $\lim_{N\to\infty} ||\mathsf{e}[\widehat{\sigma}] - \sigma^*|| = \lim_{N\to\infty} ||\mathsf{e}[\sigma_{\mathcal{M}}] - \sigma^*|| = 0.$

(iii) $\lim_{N\to\infty} \mathbb{E}_{\widehat{\sigma}}[T_{\widehat{\tau}}] = \lim_{N\to\infty} \min_{\tau\in\mathcal{M}} \mathbb{E}_{\sigma_{\mathcal{M}}}[T_\tau] = \min_{\tau\in\mathcal{M}} \mathbb{E}_{\sigma^*}[T_\tau].$

This implies Proposition 1 since for any consumption strategy σ, profits $\Pi(\theta)$ are a function of Markov average $\mathsf{e}[\sigma]$ and expected transfers $\mathbb{E}_\sigma[T]$.

Proof heuristic: Focus on the generous contract $\widehat{\tau}$ and the associated consumption strategy $\widehat{\sigma}$. The main step of the proof establishes that even under contract $\widehat{\tau}$, the consumer does not benefit from using a dynamic strategy such that total transfers $\widehat{T}$ do not converge in probability to their expectation. Dependency of expectations on type θ is suppressed for concision.

By definition of $\widehat{T}$, we have that

$$\max_{\sigma\in\Sigma} \mathbb{E}_\sigma \left[U - \widehat{T}\right] \geq \max_{\sigma\in\Sigma} \left(\mathbb{E}_\sigma[U] - \min_{\phi\in\Phi} \mathbb{E}_\sigma[T_\phi]\right). \tag{4}$$

Taking $\widehat{\sigma}$ as given, the no-regret learning literature (Blackwell, 1956, Hannan, 1957) implies that there exists a dynamic contract choice strategy ϕ such that

$$\mathbb{E}_{\widehat{\sigma}}[|\widehat{T} - T_\phi|] = O(1/\sqrt{N}).$$ [13]

Hence, in both the single and dynamic choice paradigms, the following upper bound holds,

[13] In the single choice context, this strategy need not be available to the consumer. It is just a proof device.

$$\begin{aligned}\mathbb{E}_{\widehat{\sigma}}[U - \widehat{T}] &\leq \mathbb{E}_{\phi,\widehat{\sigma}}[U - T_\phi] + O(1/\sqrt{N}) \\ &= \mathbb{E}_{\mathsf{e}[\phi,\widehat{\sigma}]}[U - T_\phi] + O(1/\sqrt{N}).\end{aligned}$$

Given Assumption 3, this implies that $e[\phi, \widehat{\sigma}]$ is an approximate solution to

$$\max_{\nu \in \Delta(\mathcal{M} \times Q^{A \times K})} \mathbb{E}_{q,\alpha|\nu} \left[U - T_\tau\right].$$

Since this program admits a unique solution, this implies points (i), (ii) and (iii) for $\widehat{\sigma}$. ■

5 The General Case

In real life setting, the distribution of preferences frequently change over time. In the retail electricity market, seasonality, life cycle effects, and housing changes all affect individual consumption preferences. In the insurance market, age, and chronic diagnostics have persistent effect over preferences.

This section extends the analysis to more general processes for the underlying private preferences α_t and public state k_t. Proposition 1 does not extend as is under arbitrarily general dynamic environments.

This is clearly the case under the repeated choice paradigm. If the distribution of α_t and k_t changes over time, then repeated choice from menus allows the consumer to tailor contract choice to temporary circumstances. In principle this allows the consumer to achieve greater utility than even generous long-term contracting in which the same contract applies to the entire period, even if the contract is specified at the end.

Still, this section is able to establish the following results:

- In "rapid learning" environments, where consumers are quickly able to predict the best contract ex post with high precision, then under the single choice paradigm, menus and generous long-term contracts are approximately equivalent.

- In "slow learning" environments, where consumers are unable to quickly predict the best contract ex post, then for a natural class of alignment-enhancing contracts, under either the single or repeated choice paradigm, menus and generous long-term contracts guarantee approximately the same lower bound on Pareto efficiency gains. However, menus and generous long-term contracts do not necessarily induce approximately similar outcomes beyond this lower bound.

5.1 Rapid learning

This section extends Proposition 1 to environments in which the consumer is able to confidently predict which contract will benefit them most. Before formalizing this idea, it is useful to clarify that optimal consumption rules given a contract are prior independent. The prior only matters to select the optimal contract.

For simplicity, let the following assumption hold.

Assumption 4. *Assume that for all $\tau \in \mathcal{M}$, and all $(\alpha, k) \in A \times K$,*

$$\max_{q \in Q} u(q, \alpha) - \tau(q, k)$$

has a unique solution.

This assumption is generic since Q, A and K are finite.

Lemma 4 (Optimal consumption is Markov)**.** *Given a transfer scheme τ, the solution to*

$$\max_{\sigma \in \Sigma} \mathbb{E}\left[\frac{1}{N}\sum_{t=1}^{N} u(q_t, \alpha_t) - \tau(q_t, k_t)\bigg|\theta\right]$$

is independent of θ and of the distribution of preferences and states $(\alpha_t, k_t)_{t\in\{1,\dots,N\}}$. Specifically, $\sigma \in Q^{A\times K}$, and for all α, k,

$$\sigma(q,k) = \arg\max_{q\in Q} u(q,\alpha) - \tau(q,k).$$

Let σ_τ denote the corresponding Markov strategy.

Proof. This is simply the policy that maximizes the integrand point by by point. □

Rapid learning is formalized as follows. Let $\mu_N \in \Delta(A\times K)$ denote the realized sample joint distribution of preferences and states (α, k) defined by

$$\mu_N(\alpha,k) \equiv \frac{1}{N}\sum_{t=1}^{N} \mathbf{1}_{(\alpha_t,k_t)=(\alpha,k)}$$

Assumption 5. *As N grows large, the consumer is approximately certain of the contract that will benefit them the most under the associated optimal consumption policy:*

$$\lim_{N\to\infty} \textit{prob}\left(\mu_N \text{ s.t. } \arg\max_{\tau\in\mathcal{M}} \mathbb{E}_{\mu_N}[U - T_\tau|\sigma_\tau] = \tau^*\right) = 1.$$

Assumption 5 holds in the conditionally i.i.d. environment of Section 4. Note that there is no actual learning happening here: the consumer has a very confident prior in the first period. Assumption 5 should be thought of as holding because learning has already occurred rapidly. In environments where the consumer is initially uncertain about their type but learns rapidly, Assumption 5 would hold at a time t_N growing sub-linearly with N.

As in Section 4, let $(W_{\mathcal{M}}(\theta), \Pi_{\mathcal{M}})_{\theta\in\Theta}$ and $(\widehat{W}(\theta), \widehat{\Pi}(\theta))_{\theta\in\Theta}$ denote expected consumer welfare and producer profits associated with each type θ given a neutral stance, respectively under menu $\mathcal{M} = \{\tau^0, \tau^1\}$ and generous long-term contract $\widehat{\tau}$.

Proposition 2. *Under the single choice paradigm, for all $\theta \in \Theta$, the following hold:*

(i) $\widehat{W}(\theta) \geq W_{\mathcal{M}}(\theta)$;

(ii) for any $\epsilon > 0$, for N large enough, $\widehat{\Pi}(\theta) \geq \Pi_{\mathcal{M}}(\theta) - \epsilon$

5.2 Slow learning

Menus and generous contracts can induce different outcomes. Propositions 1 and 2 do not extend as is in environments where Assumption 5 does not hold, even under the single choice paradigm.

As an illustration, return to the toy model of Section 2. Assume that for periods $\rho N + 1$ to N, with $\rho \in (0, 1)$, the types and the distribution of preferences and costs of production are exactly as in Section 2, with types θ_{indep} having prior probability p. For the first ρN periods, the marginal cost of production k is equal to $\underline{k}$, and the consumer's preference α is equal to a fixed value α_0 such that,

$$\underline{k} \leq \alpha_0 \leq (1-p)k_0 + p\underline{k}.$$

Assume that the consumer gets no signal of their type until period $\rho N + 1$, at which point they learn their type. Assume also that p is high enough so that under choice from menu $\mathcal{M}$ the uninformed consumer would pick variable price contract τ^1. Then it follows that under choice from menu $\mathcal{M}$ the consumer picks contract τ^1 and consumes $q = q_0$ in the first ρN periods. In contrast, under generous long-term contract $\widehat{T}$, the consumer applies a shadow price of $(1-p)k_0 + p\underline{k}$ to its consumption in the first ρN periods. As a result, the consumer chooses to consume $q = 0$ in the first ρN periods.

The case of Pareto alignments. While generous long-term contracts need not approximate the performance of reference menus in non-stationary environments, it is possible to show that they can achieve a meaningful share of Pareto efficiency gains for a specific class of reference menus.

Take as given a baseline contract $\tau^0(q, k)$, and let $\sigma_0 \in Q^{A \times K}$ denote the optimal consumption policy under this contract. Let $\pi^0(q,k) \equiv \tau^0(q,k) - c(q,k)$ denote producer flow profits, and let $\overline{\pi}^0(k)$ denote an upper bound to expected profit estimates, so that

$$\forall\theta, \forall h_t, \quad \overline{\pi}^0(k) \geq \mathbb{E}\left[\pi^0(q,k)|\theta,k,\sigma_0,h_t\right]. \tag{5}$$

Note that $\overline{\pi}^0(k)$ may have to be a crude bound but can be allowed to depend on exogenous characteristics of the consumer visible to the producer via public state k. Since it is a bound on profits under the baseline contract, the producer will have access to significant data regarding behavior under that contract, making it plausible that the producer can in fact compute fairly tight upper bounds to conditionally expected profits. It is important that public state k be exogenous to the consumption problem, i.e. that consumer choices do not change k.

An alternative contract τ^1 is a Pareto alignment of contract τ^0 if it takes the form

$$\tau^1(q,k) = \tau^0(q,k) - \rho \times (\pi^0(q,k) - \overline{\pi}^0(k))$$

for $\rho \in (0,1)$.

Let $v^0(q,k,\alpha)$, $v^1(q,k,\alpha)$, and $\pi^0(q,k)$, $\pi^1(q,k)$ respectively denote flow payoffs to the consumer and the producer under contracts τ^0 and τ^1. Contracts τ^1 induce flow payoffs for the producer and the consumer taking the following form:

$$\begin{aligned}
\pi^1(q,k) &= \pi^0(q,k) - \rho \times [\pi^0(q,k) - \overline{\pi}^0(k)] \\
&= \overline{\pi}^0(k) + (1-\rho) \times [\pi^0(q,k) - \overline{\pi}^0(k)] \qquad (6) \\
v^1(q,k,\alpha) &= v^0(q,k,\alpha) + \rho \times [\pi^0(q,k) - \overline{\pi}^0(k)] \\
&= v^0(q,k,\alpha) + \frac{\rho}{1-\rho} \times [\pi^1(q,k) - \overline{\pi}^0(k)] \qquad (7)
\end{aligned}$$

In other terms, contract τ^1 induces interim preferences for the consumer that place additional weight $\rho/(1-\rho)$ on the payoff to the principal.

It follows from (6) and (7) that offering menu $\mathcal{M}$ necessarily leads to a weak Pareto improvement compared to simply offering baseline contract τ^0.

Let $(W_0(\theta), \Pi_0(\theta))_{\theta\in\Theta}$, $(W_{\mathcal{M}}(\theta), \Pi_{\mathcal{M}}(\theta))_{\theta\in\Theta}$ and $(\widehat{W}(\theta), \widehat{\Pi}(\theta))_{\theta\in\Theta}$ denote expected consumer welfare and producer profits associated with each type θ, given a neutral stance, respectively under contract τ^0, menu $\mathcal{M} = \{\tau^0, \tau^1\}$ and generous long-term contract $\widehat{\tau}$.

Proposition 3 (Pareto alignments)**.** *Under the neutral stance, in both the single and repeated choice paradigms, and for every type θ, menu $\mathcal{M}$ leads to a weak Pareto improvement over contract τ^0.*

Furthermore,

$$\Pi_{\mathcal{M}}(\theta) \geq \Pi_0(\theta) + \frac{1-\rho}{\rho}\left[W_{\mathcal{M}}(\theta) - W_0(\theta)\right]. \tag{8}$$

Proposition 3 relies on k being exogenous to the consumer's consumption choices, but does not depend on condition (5).

Payoff bound (8) shows up in robust contracting work studying linear contracts (Chassang, 2013, Carroll, 2015) and linear alignments (Madarász and Prat, 2017). It is made tight by environments in which global IC constraints were near-binding under τ^0, i.e. the consumer was approximately indifferent between their chosen consumption plan, and a consumption plan that benefited the producer. Indeed, typically, increasing the producer's payoff would require making decisions that reduce the consumer's baseline utility. This effect is minimized when the consumer was initially near indifferent with a consumption strategy that made the principal much better off. Contract τ^1 induces the consumer to change their consumption to the producer's preferred consumption plan. This maximize the consumer's benefit associated with the producer's increased profit, leading to a tight lower bound.

It turns out that payoff bound (8) holds approximately under the generous contract $\widehat{\tau}$ derived from menu $\mathcal{M}$. This result relies on condition (5).

Proposition 4. *(i) For all θ, contract $\widehat{\tau}$, weakly improves the consumer's welfare over τ: $\widehat{W}(\theta) - W_0(\theta) \geq 0$.*

(ii) There exists a fixed $M > 0$ such that for any N,

$$\widehat{\Pi}(\theta) \geq \Pi_0(\theta) + \frac{1-\rho}{\rho}\left[\widehat{W}(\theta) - W_0(\theta)\right] - P_N,$$

with

$$P_N \equiv \mathbb{E}\left(\left[\sum_{t=1}^{N} \pi^0(q_t^0, k_t) - \overline{\pi}^0(k_t)\right]^+\right) \leq \frac{M}{\sqrt{N}}. \tag{9}$$

In words, generous long-term contracts achieve the same performance guarantee as Pareto alignments, up to a penalty that vanishes as the contracting horizon becomes large.

6 Discussion

This section concludes the paper with rough calibrations in two contexts, load-shedding and healthcare insurance, as well as a discussion of practical implementation concerns.

6.1 Calibration in the context of load-shedding

Proposition 4 lends itself to straightforward calibration, since the magnitude of penalty term P_N can be evaluated using only data from the consumption process under the baseline contract τ^0. A key design aspect is the upper-bound $\overline{\pi}$ on conditionally expected profits used for Pareto alignment.

A natural strategy is to apply an inflation coefficient $c > 0$ to an estimate of benchmark profits conditional on available information:

$$\overline{\pi}^0(k) = (1+c)\mathbb{E}[\pi^0(q,k)|k].$$

For c large enough, condition (5) holds, but that is not sufficient for Pareto alignments to be effective. If c is too high, then the alternative contract τ^1 is never ex post optimal, and hence irrelevant to behavior: both penalty P_N and the impact of offering τ^1 on behavior are small. Inversely, if c is too small, then penalty P_N is no longer small relative to potential profit increases.

The trade-off between potential efficiency gains and profit penalties is mediated by the quality of exogenous information used to formulate baseline profit expectations $\mathbb{E}[\pi^0|k]$. Better conditioning allows to keep coefficient c small, while keeping penalty P_N small, reducing the intensity of the trade-off. The difficulty is that in principle, the conditioning should be based on exogenous variables, so as not to affect incentives in the first place. In practice, utilities can collect relevant characteristics that help predict consumption patterns, for instance, whether the consumer uses electric heating or not. In addition, it may be possible to use past consumer-specific profits to estimate current baseline profits without creating significant moral hazard. As Figure 1 shows, in the French context the cost of electricity is frequently above and below base contract prices, especially in periods where the cost of electricity is high. This makes it difficult for consumers to figure out how changes in their demand would affect the producer's profits on an average day. For this reason, it may be possible to use some information about past individual consumption without significantly changing consumption behavior.

Data. The calibration is performed using data associated with the French context described in Section 2. The goal is to get a sense of order of magnitudes.

Industry reports suggest that consumers joining variable price contracts (referred to as Tempo contracts) can save between 20% and 30% on their electricity bill. The current calibration will focus on a Pareto alignment of the baseline fixed price contract with surplus sharing parameter $\rho = 1/2$. To fix ideas, assume that this Pareto alignment leads to an overall surplus improvement equal to 20% of profits, that is allocated equal parts between the consumer and producer.

Fixed price contract details as well as electricity market spot prices are publicly available from grid operator RTE (RTE, 2025). Temperature data, an important predictor of consumption is also publicly available, and so is aggregate electricity consumption data. The difficulty is to obtain individual consumption data, and individual covariates that potentially predict persistent differences in the profitability of consumers. Such data consitutes personally identifying information under European privacy regulation (known as RGPD), and is not readily available. However, legacy operators RTE and EDF have made available datasets of simulated individual temperature and consumption patterns exploiting large neural networks to create highly realistic consumption series (Nabil et al., 2025). The data consists of a year of consumption choice for 10 000 hypothetical households (with daily total consumption closely matching actual consumption in 2023) and local temperature at the half-hour time step. Individual demand and spot prices are combined to obtain individual profitability at the hourly level, treating spot prices as the cost of electricity.

It is instructive to study persistent individual differences in consumption patterns. Table 1 reports coefficient estimates and R^2 for three models of predicted consumption:

$$Q_{i,t} = \text{Pooled HDM FE} + \beta_{HDD}[15 - D_{i,t}]^+ + \beta_{CDD}[D_{i,t} - 20]^+ + \text{i.i.d. errors}$$
$$Q_{i,t} = \text{Individual HDM FE} + \beta_{HDD}[15 - D_{i,t}]^+ + \beta_{CDD}[D_{i,t} - 20]^+ + \text{i.i.d. errors}$$
$$Q_{i,t} = \text{Individual HDM FE} + \beta_{HDD}[15 - D_{i,t}]^+ + \beta_{CDD}[D_{i,t} - 20]^+ + \text{AR(1, 24) errors}$$

where Q denotes hourly consumption in kWh, HDM denotes hourly, day-of-week, and month fixed effects, D is the hourly temperature in degrees Celsius, HDD and CDD respectively refer to heating and cooling degree days, while $AR(1, 24)$ refers to error processes auto-correlated at the hourly and daily level.[14] The dynamic adjusted R^2 uses realized residuals to predict consumption and corrects for the number of degrees of freedom. Controlling for lagged residuals only makes a difference when the model of errors allows for auto-correlation.

[14] Coefficients on the 1h and 24h lags are 0.25 and 0.45. Including a 24 hour lag yields a large improvement in R^2.

	Pooled FE + i.i.d. err	Indiv FE + i.i.d. err	Indiv FE + AR err
β_{CDD}	0.0258	0.0303	0.0303
β_{HDD}	0.0035	0.0025	0.0025
Dyn. Adj. R^2	0.10	0.53	0.67

Table 1: Persistent differences in individual consumption behavior are large

From the perspective of keeping penalty P_N small, the main takeaway is that the bulk of the variation in consumption is explained by persistent individual level heterogeneity, and that much of the heterogeneity is associated with persistent shocks rather than individual fixed effects. This suggests that there will be considerable value in controlling for recent individual consumption when predicting individual counterfactual profitability.

Findings. The calibration contrasts two conditioning strategies used to estimate counterfactual reference profits $\overline{\pi}^0(k_t)$ that likely bracket the magnitude of penalties plausibly achievable in practice.

The first strategy simply consists in applying a fixed inflation coefficient to the empirical population mean at each date t. For $c > 0$,

$$\overline{\pi}^0_{i,t} \equiv (1+c) \times \widehat{\mathbb{E}}_{OLS}\left[\pi_{i,t} | \mathsf{temp}_{i,t}\right]$$

where

$$\widehat{\mathbb{E}}_{OLS}\left[\pi_{i,t} | \mathrm{temp}_{i,t}\right]$$

is the OLS estimate of profit $\pi_{i,t}$ given local temperature $\mathsf{temp}_{i,t}$ computed using the population level data at date t.

This strategy uses only uses local temperature to predict individual consumption, which makes it poorly suited to capture persistent differences in consumption across individuals. This implies that a fairly high inflation coefficient c is needed to keep penalties small. The top panel of Figure 2 (penalty for pooled profit estimate) illustrates the ratio of penalty to baseline aggregate profits, P_N/Π^0, for various inflation coefficients c. For $c = 0$, the penalty is greater than a potential profit upside of 10%. For significantly larger inflation coefficients (e.g. 50%) the penalty falls to 2% or below, but the potential profit upside is also likely to be much smaller since it becomes unlikely that the Pareto aligned contract τ^1 is the ex post optimal one for the consumer.

The second strategy, illustrated by the bottom panel of Figure 2 (penalty for individual profit estimate), explicitly targets persistent individual differences in consumption. Specifically, it computes an individual time-varying relative profit ratio $\lambda_{i,t}$, by taking an expanding mean of the absolute value of individual profits $\pi^0_{i,s}$ divided by the expanding mean of the absolute value of population profits π^0_s. This relative profit ratio is then multiplied with contemporaneous population profits:

$$\overline{\pi}^0_{i,t} \equiv (1+c)\lambda_{i,t}\pi^0_t \quad \text{with} \quad \lambda_{i,t} \equiv \frac{\sum_{s<t} |\pi^0_{i,s}|}{\sum_{s<t} |\pi^0_s|}.$$

Because this strategy accounts for persistent individual differences in profitability, it yields much lower penalties for much lower inflation coefficients. The relative penalty is 5% for $c = 0$ and falls well under 1% for $c = 10\%$.

The takeaway is that because of persistent individual differences in profitability, it is essential to formulate individualized estimates of counterfactual profitability to ensure that generous Pareto aligned contracts are effective. Ideally, relative profitability consumption coefficient $\lambda_{i,t}$ would be assigned on the basis of observed persistent characteristics (like heating technology, housing type. . .) rather than past consumption under the baseline contract. In that sense, the performance of the individualized counterfactual estimate is likely an upper bound to the performance one could expect from generous Pareto alignments. Still, because it is tricky to predict whether increasing consumption increases or decreases profits, using averaged-out profitability data need not create significant moral hazard.

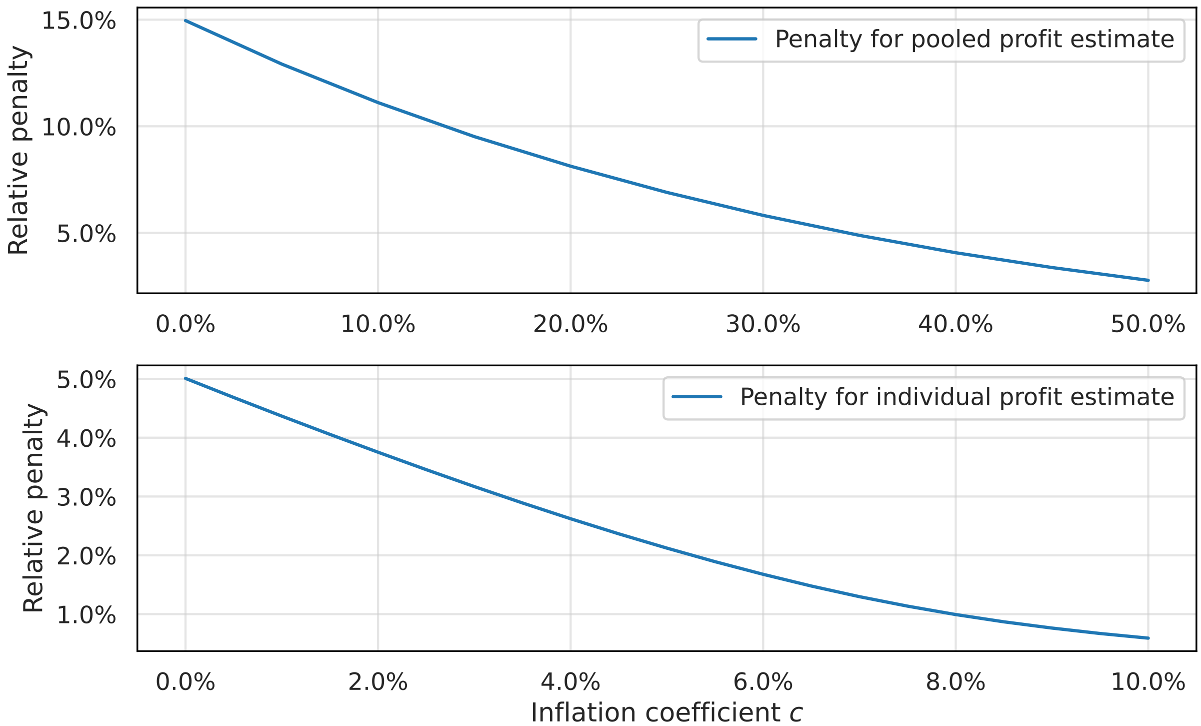

Figure 2: Penalty relative to profits by conditioning strategy and inflation coefficient

6.2 Calibration in the context of healthcare insurance

Public insurers in Europe and large employers in the USA are often mandated or under pressure to offer insurance contracts with low copays on medical visits, as well as drugs. This can result in high and potentially inefficient spending. Consistent with large scale experimentations (Newhouse and Insurance Experiment Group, 1993, Finkelstein et al., 2012) showing that health care demand is sensitive to marginal costs, there has been a push to introduce low-base-cost high-deductible health plans (HDHP) as an alternative to high-base-cost low-copay low-deductible contracts (associated in the US with preferred provider organizations, or PPOs).

Brot-Goldberg et al. (2017) provides extensive details on the coexistence of PPOs and HDHPs as well as on the impact of removing PPOs from menus in the context of a plan change at a large US firm. When both contracts are offered, 85% of employees choose the low deductible low copay PPO. This is consistent with the idea that demand for high powered, less predictable contracts tends to be weak. Removing the PPO option reduces spending by roughly 10 to 15%, but the impact is equally spread around high and low medical value spending (e.g. preventive care vs. imaging). This suggests that PPOs may in fact be desirable contracts for some subset of patients, making it difficult for most firms and policy makers to get rid of them entirely.

Altogether, this context matches the tension motivating the current paper. There is potential value to the introduction of contracts that align incentives better than PPOs, but it is difficult to do so in a manner that attracts significant demand while remaining profitable. This lead the firm to drop the PPO contract altogether. Generous long-term contracts offer an alternative.

Low deductible, low copay contracts correspond to a benchmark $\tau_0(q,k) = p_0$ with p_0 a fixed premium independent of the quantity of care consumed, and k an observable set of patient diagnostics. As in Section 5, the alternative contract τ_1 is a Pareto alignment of τ_0, taking the following form

$$\tau^1(q,k) = p_0 - \rho \times (\overline{c}^0(k) - c(q,k)) \quad \text{for} \quad \rho \in (0,1),$$ [15]

where term $\overline{c}^0(k) - c(q,k)$ captures the reduction in health care costs compared to benchmark expectations under contract τ_0. Condition (5) can be reexpressed as

$$\overline{c}^0(k) \leq \mathbb{E}[c^0(q,k)|\theta, k, \sigma_0, h_t].$$

[15]Under this family of Pareto alignments, patients are incentivized to reduce their health spending at a constant rate, off of a benchmark cost estimate dependent on exogenous diagnostics. The fact that incentives for cost reduction are stable through time differentiates them from HDHPs which provides incentives only to patients anticipating low medical cost at the yearly level. This is partly because HDHPs seek to both provide incentives and offer some risk protection. In the menu of contracts $\{\tau_0, \tau_1\}$, the presence of full insurance contracts τ_0 offers cost protection, while τ_1 offers consumers capable of reducing usage the incentive to do so. Anchoring cost estimates on a counterfactual also absorbs significant risk even when patients are under higher powered incentives. In principle, this is an important fix since Brot-Goldberg et al. (2017) find that almost all of the savings associated with HDHPs happen in months where patients' spending is below the deductible, and half of reductions are associated with patients who should expect high costs.

As in the case of load-shedding, a natural strategy is to consider cost predictions from the class of deflated coarsened cost estimates

$$\overline{c}^0(k) \equiv (1-\delta)\mathbb{E}[c^0(q,k)|\theta,k,\sigma_0]$$

with $\delta \in (0,1)$.[16]

Data. As in the case of load-shedding penalty P_N is calibrated using synthetic medical cost data made available by the Centers for Medicare & Medicaid Services (2014). The data consists of a yearly panel of diagnostics, and expenses for a representative set of medicare beneficiaries for the years 2008, 2009, and 2010. Cost estimate $\mathbb{E}[c^0(q,k)|\theta,k,\sigma_0]$ is built using random forest regressors applied to available diagnostics, with mean-absolute-error as the minimization objective, consistent with the fact that penalty P_N is proportional to the positive part of prediction errors. Penalties are computed for a 3 year contract. In the context of medicare, since no competing contract exists, it would be plausible to consider longer-term contracts. Taking 80 as a conservative estimate of the life expectancy of US 65 year-olds, a 15 year contract duration may be plausible. Given the low auto-correlation in errors, this would shrink penalties by roughly 65%.

Findings. Figure 3 summarizes findings, setting $\rho = 1/2$, and varying the cost deflation coefficient δ. In words, for $\delta = .1$, a consumer benefits from contract τ_1 if they can bring their realized costs 10% below the estimated costs of patients with similar diagnostics. To set expectations over potential cost reductions, note that Brot-Goldberg et al. (2017) find an overall impact of HDHPs on spending on the order of 15%.[17]

[16] Brot-Goldberg et al. (2017) emphasizes the importance of not driving out preventive care. In principle, expenses that reduce future spending can be captured by setting $c^0(q,k)$ as an estimate of the change in the net present value of costs rather than flow accounting costs. Appendix A.2 of Braverman and Chassang (2022) provides of formal description of this point in the context of medicare advantage.

[17] The range of effects goes roughly from 10% to 25% depending on sample selection and comparison group.

For deflation coefficients $\delta \in \{.05, .1, .15\}$, penalties are on the order of 6%, 5% and 4% for the 3 year contract, and 2.7%, 2.25% and 1.8% for the 15 year contract. While the penalties are significant they suggest that generous dynamic contract offers a plausible strategy to encourage the adoption of better aligned insurance contracts.

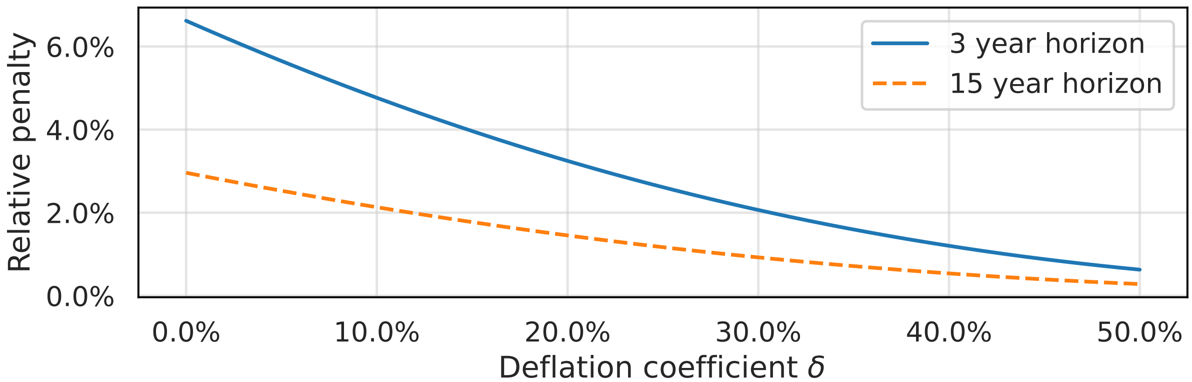

Figure 3: Penalty relative to expenditures by cost deflation coefficient

6.3 Practical Considerations

Envisaging the concrete implementation of generous long-term contracts raises addition challenges discussed below.

Capital investments. In the load-shedding context, consumers that sign up for variable price contracts sometimes make complementary capital investments helping them modulate their use. For instance, consumers using electric heating will invest in a wood fired stove to manage high electricity-cost days in the winter.

Although the framework involves repeated small consumption decisions, Propositions 1, 2, 3, and 4 extend as is to environments in which some rare actions have large payoff consequences, such as choosing whether to install a gas furnace, or an electric heat pump. Recall however that under non-stationary environments with slow learning, generous Pareto aligned contracts replicate only the payoff guarantee of the associated menu. However, if the menu achieves a higher payoff to the principal than the lower bound, that payoff need not be achieved by the generous Pareto aligned benchmark, especially under the single decision benchmark.

Leaving and returning. Consumer protection law often limit producers' ability to prevent consumers from leaving utility contracts. This is the case in the French retail electricity market, where consumers cannot be locked in with a producer for the long-term: they always retain the penalty-free ability to move to a different electricity provider. However, this does not mean that they can later return to their original provider with access to the same contractual offer as completely new clients. The producer is allowed to use cost-relevant historical data to specify contractual terms for returning customers.

All our results continue to hold in an environment where consumers can leave and return to a provider if they are treated according to their entire history with the provider. In other terms, periods t and $t+1$ need not refer to contiguous periods, but rather successive periods of interaction with the same utility provider.

Simplicity and user interface. Jessoe and Rapson (2014) highlight the importance of providing consumers with relevant information, and more generally building a convenient consumer interface in order to increase the elasticity of consumption to prices.

From this perspective, generous long-term contracting potentially adds an unwelcome dimension of complexity: correct optimization requires anticipating the likelihood that one contract or the other will be relevant. In the context of health insurance contracts with high deductibles, Brot-Goldberg et al. (2017) show that consumers facing non-linear contracts often do not correctly anticipate the correct shadow cost of care. Patients respond to spot prices at current health spending rather than the true expected shadow cost of care.

On the whole, the argument for keeping consumer interfaces simple, primarily informing them of the sensitivity of flow contract $\widehat{\tau}$ to consumption, is persuasive. While this may lead to incorrect decision making early on, as the ex post optimal contract becomes clearer over time, this will provide accurate information for decision-making. This is especially true for open-ended long-term generous contracts: over time, the spot price matches the shadow price, and the cost of erroneous decision-making has a chance to mechanically vanish. This is more difficult under HDHPs whose spot price resets on a yearly basis.

In the spirit of Tempo contracts which group days in few categories, it may be useful to coarsen the incentives provided to consumers. Concretely Pareto alignments need not offer the same profit sharing rule every day, and it may be preferable to set binary priorities, in which profit sharing is high on days where alignment matters most, and set to 0 in other periods. This may simplify the attention problem consumers need to solve.

Appendix

A Proofs

A.1 Proofs for Section 4

Results for this section hold in the more general case where preferences take the form

$$\frac{1}{N}\sum_{t=1}^{N} u(q_t, \alpha_t) - \Gamma\left(\frac{1}{N}\sum_{t=1}^{N} \tau(q_t, k_t)\right),$$

where Γ is increasing, convex, and Lipschitz continuous with modulus $L > 0$.

Assumption 3 is generalized as follows.

Assumption A.1. *For every type θ, mapping*

$$\nu \in \Delta\left(\mathcal{M} \times Q^{A\times K}\right) \mapsto \mathbb{E}_{q,\alpha|\nu,\theta}\left[U\right] - \Gamma\left(\mathbb{E}_{q,k|\nu,\theta}\left[T_\tau\right]\right) \tag{A.1}$$

has a unique maximizer ν^*_θ, and it is in pure strategy with respect to $\tau \in \mathcal{M}$, i.e. $\nu^*_\theta = (\tau^*_\theta, \sigma^*_\theta) \in \mathcal{M} \times \Delta(Q^{A\times K})$.

Note that the assumption that unique maximizer ν^*_θ is in pure strategies with respect to τ is without loss of generality for Assumption 3, but not here, since convex cost Γ provides a motive for using mixtures.

Proposition 1 extends as is, and Lemma 3 extends as follows.

Lemma A.1 (strategy averages converge)**.** *Under either the repeated or single choice paradigm, as N grows large, the following hold*

(i) Markov averages $\mathsf{e}[\widehat{\sigma}]$ and $\mathsf{e}[\sigma_{\mathcal{M}}]$ asymptotically solve

$$\max_{\sigma\in\Delta(Q^{A\times K})} \mathbb{E}_{q,\alpha|\sigma,\theta}\left[U\right] - \min_{\tau\in\mathcal{M}} \Gamma\left(\mathbb{E}_{q,k|\sigma}\left[T_\tau\right]\right).$$

(ii) $\lim_{N\to\infty} ||\mathsf{e}[\widehat{\sigma}] - \sigma^*|| = \lim_{N\to\infty} ||\mathsf{e}[\sigma_{\mathcal{M}}] - \sigma^*|| = 0.$

(iii) $\lim_{N\to\infty} \mathbb{E}_{\widehat{\sigma}}[T_{\widehat{\tau}}] = \lim_{N\to\infty} \min_{\tau\in\mathcal{M}} \mathbb{E}_{\sigma_{\mathcal{M}}}[T_\tau] = \min_{\tau\in\mathcal{M}} \mathbb{E}_{\sigma^*}[T_\tau].$

The following result is helpful before turning to the proof of Lemma A.1

Lemma A.2. *The dependency of expectations on θ is suppressed for the sake of simplicity. Let $\sigma \in \Delta(Q^{A\times K})$ denote a Markov perfect consumption strategy. The following hold uniformly over σ:*

$$\forall \tau \in \mathcal{M}, \quad \mathbb{E}_\sigma\left[U - \Gamma(T_\tau)\right] = \mathbb{E}_\sigma\left[U\right] - \Gamma\left(\mathbb{E}_\sigma\left[T_\tau\right]\right) + O(1/\sqrt{N}) \tag{A.2}$$

$$\mathbb{E}_\sigma\left[U - \Gamma(\widehat{T})\right] = \mathbb{E}_\sigma\left[U\right] - \min_{\tau\in\mathcal{M}} \Gamma\left(\mathbb{E}_\sigma\left[T_\tau\right]\right) + O(1/\sqrt{N}) \tag{A.3}$$

In words, for Markov consumption strategies, it is possible to replace realized transfers with their expectation.

Proof. Let

$$\Delta_\tau \equiv \frac{1}{N}\sum_{t=1}^{N} \tau(h_t) - \mathbb{E}\left[\frac{1}{N}\sum_{t=1}^{N} \tau(h_t)\right].$$

Δ_τ is an average of bounded martingale increments. Applying the Azuma-Hoeffding theorem to $|\Delta_\tau|$, and using the fact that Γ is Lipschitz with coefficient and it follows that there exists a constant M independent of σ such that

$$\begin{aligned} & \mathbb{E}_\sigma\left[\Gamma(\mathbb{E}_\sigma\left[T_\tau\right]) - L|\Delta_\tau|\right] \le \mathbb{E}_\sigma\left[\Gamma(T_\tau)\right] \le \mathbb{E}_\sigma\left[\Gamma(\mathbb{E}_\sigma\left[T_\tau\right]) + L|\Delta_\tau|\right] \\ \Rightarrow \quad & \Gamma(\mathbb{E}_\sigma\left[T_\tau\right]) - ML\frac{1}{\sqrt{N}} \le \mathbb{E}_\sigma\left[\Gamma(T_\tau)\right] \le \Gamma(\mathbb{E}_\sigma\left[T_\tau\right]) + ML\frac{1}{\sqrt{N}}. \end{aligned}$$

This establishes (A.2).

Turn to (A.3). The following upper bound holds:

$$\mathbb{E}_\sigma\left[\Gamma(\widehat{T})\right] \le \min_{\tau\in\mathcal{M}} \mathbb{E}_\sigma[\Gamma(T_\tau)] \le \min_{\tau\in\mathcal{M}} \Gamma(\mathbb{E}_\sigma\left[T_\tau\right]) + ML\frac{1}{\sqrt{N}}.$$

For a matching lower bound, observe that

$$\widehat{T} \geq \min_{\tau \in \mathcal{M}} \mathbb{E}\left[T_\tau\right] + \min_{\tau \in \mathcal{M}} \Delta_\tau \Rightarrow \widehat{T} \geq \min_{\tau \in \mathcal{M}} \mathbb{E}\left[T_\tau\right] - \sum_{\tau \in \mathcal{M}} |\Delta_\tau|,$$
$$\Rightarrow \Gamma(\widehat{T}) \geq \min_{\tau \in \mathcal{M}} \mathbb{E}_\sigma\left[\Gamma(T_\tau)\right] - L \sum_{\tau \in \mathcal{M}} |\Delta_\tau|.$$

It follows from applying the Azuma-Hoeffding inequality that there exists M independent of σ such that

$$\mathbb{E}_\sigma\left[\Gamma(\widehat{T})\right] \geq \min_{\tau \in \mathcal{M}} \mathbb{E}\left[\Gamma(T_\tau)\right] - 2ML\frac{1}{\sqrt{N}}.$$

This concludes the proof. □

Proof of Lemma A.1. Consider the case of $\widehat{\sigma}$. Recall that, suppressing dependency on type θ, σ^* solves

$$\max_{\sigma \in \Delta(Q^{A \times K})} \mathbb{E}_{q,\alpha|\sigma}\left[U\right] - \min_{\tau \in \mathcal{M}} \Gamma\left(\mathbb{E}_{q,k|\sigma}\left[T_\tau\right]\right).$$

By definition of $\widehat{\sigma}$,

$$\mathbb{E}_{\widehat{\sigma}}[U - \Gamma(\widehat{T})] \geq \mathbb{E}_{\sigma^*}[U - \Gamma(\widehat{T})].$$

Using Lemma A.2, this implies that

$$\mathbb{E}_{\widehat{\sigma}}[U - \Gamma(\widehat{T})] \geq \mathbb{E}_{\sigma^*}[U] - \min_{\tau \in \mathcal{M}} \mathbb{E}_{\sigma^*}[\Gamma(T_\tau)] - O(1/\sqrt{N}).$$

Taking $\widehat{\sigma}$ as given, the no-regret learning literature (Blackwell, 1956, Hannan, 1957) implies that there exists a dynamic contract choice strategy ϕ such that

$$\mathbb{E}_{\widehat{\sigma}}[|\widehat{T} - T_\phi|] = O(1/\sqrt{N}).$$

Hence, by convexity of Γ,

$$
\begin{aligned}
\mathbb{E}_{\widehat{\sigma}}[U-\Gamma(\widehat{T})] &\leq \mathbb{E}_{\widehat{\sigma}}[U]-\Gamma\left(\mathbb{E}_{\widehat{\sigma}}[\widehat{T}]\right) \\
&\leq \mathbb{E}_{\widehat{\sigma}}[U]-\Gamma\left(\mathbb{E}_{\widehat{\sigma},\phi}[T_\phi]\right)+O(1/\sqrt{N}) \\
&= \mathbb{E}_{\mathsf{e}[\widehat{\sigma}]}[U]-\Gamma\left(\mathbb{E}_{\mathsf{e}[\phi,\widehat{\sigma}]}[T_\phi]\right)+O(1/\sqrt{N}).
\end{aligned}
$$

Given Assumption 3, this implies that $e[\phi, \widehat{\sigma}]$ is an approximate solution to

$$
\max_{\nu\in\Delta(\mathcal{M}\times Q^{A\times K})} \mathbb{E}_{q,\alpha|\nu}\left[U\right]-\Gamma\left(\mathbb{E}_{q,k|\nu}\left[T_\tau\right]\right). \tag{A.4}
$$

Since this program admits a unique solution, this implies points (i), (ii) and (iii) for $\widehat{\sigma}$ in the repeated choice paradigm.

Since the solution to (A.4) is (by Assumption A.1) in pure strategies with respect to τ, this implies that $\mathsf{e}\left[\phi\right]$ converges in probability to a point mass on a contract $\tau \in \mathcal{M}$. This implies that $\mathsf{e}\left[\widehat{\sigma}\right]$ is an approximate solution to

$$
\max_{\sigma\in\Delta(Q^{A\times K})} \mathbb{E}_{q,\alpha|\sigma}\left[U\right]-\min_{\tau\in\mathcal{M}}\Gamma\left(\mathbb{E}_{q,k|\sigma}\left[T_\tau\right]\right).
$$

This yields points (i), (ii), (iii) for $\widehat{\sigma}$ under the single choice choice paradigm.

A essentially identical proof holds in the case of σ_M. ■

A.2 Proofs for Section 5

Proof of Proposition 2. Throughout, type θ is taken as given and the dependency of expectations on θ is suppressed for readability.

Observe that although in principle, consumption could depend on past histories, solutions to

$$\max_{\sigma\in\Sigma}\max_{\tau\in\mathcal{M}}\mathbb{E}[U-T_\tau]$$

are Markovian. Indeed, given realized values of α, k, and a choice of contract τ, consumption q must solve

$$\max_{q\in Q} u(\alpha,q)-\tau(q,k)$$

which is independent of history h_t.

Let $\sigma^*\in Q^{A\times K}$ and τ^* denote the solution to $\max_{\sigma\in Q^{A\times K}}\max_{\tau\in\mathcal{M}}\mathbb{E}[U-T_\tau|\sigma]$.

The proof will show that $\widehat{\sigma}$, solution to $\max_{\sigma\in\Sigma}\mathbb{E}[U-\widehat{T}|\widehat{\sigma}]$ must coincide with σ^* at almost all histories.

Let τ_N denote the solution to the ex post optimization problem

$$\max_{\tau\in\mathcal{M}}\mathbb{E}_{\mu_N}\left[U-T_\tau|\sigma_\tau\right],$$

where σ_τ is the optimal strategy given τ which is Markov by Lemma 4. Let τ' denote the contract in $\mathcal{M}$ that is not τ^*. Finally, let $\widehat{q}_t$ denote consumption choices under $\widehat{\sigma}$ and q_t^* denote consumption choices under σ^*.

By Assumption 5,

$$\mathbb{E}_{\widehat{\sigma}}\left[U-\widehat{T}\right]\leq \mathsf{prob}(\tau_N=\tau^*)\mathbb{E}_{\widehat{\sigma}}\left[U-T_{\tau^*}|\tau_N=\tau^*\right]+o(1).$$

If $\tau_N=\tau^*$, it follows from Assumption 4 that there exists $\eta>0$ such that

$$U-\widehat{T}=\mathbb{E}_{\mu_N}\left[U-T_{\tau^*}|\widehat{\sigma}\right]\leq\mathbb{E}_{\mu_N}\left[U-T_{\tau^*}|\sigma^*\right]-\eta\mathbb{E}_{\mu_N}\left[\frac{1}{N}\sum_{t=1}^N\mathbf{1}_{\widehat{q}_t\neq q_t^*}\right].$$

Integrating over realizations of μ_N, this implies that

$$\mathbb{E}_{\widehat{\sigma}}\left[U-\widehat{T}\right]\leq\mathbb{E}_{\sigma^*}\left[U-T_{\tau^*}\right]-\eta\mathbb{E}\left[\frac{1}{N}\sum_{t=1}^N\mathbf{1}_{\widehat{q}_t\neq q_t^*}\right]+o(1)$$

Since by definition of $\widehat{\sigma}$ and $\widehat{T}$,

$$\mathbb{E}_{\widehat{\sigma}}\left[U - \widehat{T}\right] \geq \mathbb{E}_{\sigma^*}\left[U - T_{\tau^*}\right],$$

it follows that

$$\lim_{N\to\infty} \mathbb{E}\left[\frac{1}{N}\sum_{t=1}^{N} \mathbf{1}_{\widehat{q}_t \neq q_t^*}\right] = 0.$$

This implies that consumption choices and transfers under the generous contract converge to those under single contract choice from menu $\mathcal{M}$. ■

Proof of Proposition 3. For concision, the proof focuses on the repeated choice framework. An identical argument holds in the single choice framework. Since the consumer has more choice, menu $\mathcal{M}$ is necessarily a weak improvement from their perspective. For this reason, the proof focuses on the producer's profits.

For types θ choosing baseline contract τ^0 in period t, behavior and profits are unchanged.

Consider now types θ choosing contract τ^1. Let σ_0 and σ_1 in $Q^{A\times K}$ denote the respective strategies maximizing flow payoffs to the agent point by point under respective contracts τ^0 and τ^1:

$$\max_{q\in Q} u(\alpha, q) - \tau^0(q,k) \quad \text{and} \quad \max_{q\in Q} u(\alpha, q) - \tau^1(q,k)$$

Since σ_0 is optimal under contract τ^0 it follows that

$$\mathbb{E}_{\sigma_0,\alpha,k}[u(q,\alpha) - \tau^0(q,k)] \geq \mathbb{E}_{\sigma_1,\alpha,k}[u(q,\alpha) - \tau^0(q,k)] \tag{A.5}$$

Since type θ prefers contract τ^1 over τ^0, there exists $\eta \geq 0$ such that

$$\mathbb{E}_{\sigma_1,\alpha,k}[u(q,\alpha) - \tau^0(q,k) + \rho[\pi^0(q,k) - \overline{\pi}^0(k)]] = \mathbb{E}_{\sigma_0,\alpha,k}[u(q,\alpha) - \tau^0(q,k)] + \eta \tag{A.6}$$

Together (A.5) and (A.6) imply that

$$\begin{aligned}&\mathbb{E}_{\sigma_1,\alpha,k}[u(q,\alpha)-\tau^0(q,k)+\rho[\pi^0(q,k)-\overline{\pi}^0(k)]]\geq \mathbb{E}_{\sigma_1,\alpha,k}[u(q,\alpha)-\tau^0(q,k)]+\eta\\ \Rightarrow&\mathbb{E}_{\sigma_1,\alpha,k}[\pi^0(q,k)-\overline{\pi}^0(k)]\geq \frac{1}{\rho}\eta\end{aligned}$$

By (6), and since $\overline{\pi}^0(k)$ is an upper bound to expected profits under τ^0, it follows that

$$\begin{aligned}\mathbb{E}_{\sigma_1,\alpha,k}[\pi^1(q,k)]&\geq \mathbb{E}_k[\overline{\pi}^0(k)]+(1-\rho)\mathbb{E}_{\sigma_1,\alpha,k}[\pi^0(q,k)-\overline{\pi}^0(k)]\\ &\geq \mathbb{E}_{\sigma_0,\alpha,k}[\pi^0(q,k)]+\frac{1-\rho}{\rho}\eta\end{aligned}$$

This implies that the principal must always benefit when the consumer chooses τ^1, and yields bound (8) using (A.6) to express η as the consumer's increase in welfare. ■

Proof of Proposition 4. Point (i) is immediate since contract $\widehat{\tau}$ lowers the cost of consumption on a path by path basis. Turn now to point (ii).

First recall that the optimal policy σ_0 under contract τ^0 is Markov perfect: consumption choices at time t depend only on α_t, k_t.

Second, observe that total payments $T_{\widehat{\tau}}$ under the generous contract take the form

$$T_{\widehat{\tau}}=\min_{\tau\in\mathcal{M}}T_\tau=T_{\tau^0}-\rho\left[\Pi_0-\overline{\Pi}_0\right]^+.$$

Consider now optimal policy $\widehat{\sigma}$ under the generous contract $\widehat{\tau}$. Because $\widehat{\tau}$ itself depends on past histories, $\widehat{\sigma}$ is not Markov, and depends on an appropriate state variable.

Take as given the type of the consumer. For concision, dependency on type is suppressed below. Take as given a private history $h_t=(\alpha_s,k_s)_{s=\{1,\dots,t\}}\in(A\times K)$ of realized preferences and states.

For any value $\Delta\Pi\in\mathbb{R}$, and history h_t. define value function V as follows

$$V(\Delta\Pi, h_t) \equiv \max_{\sigma \in \Sigma} \mathbb{E}\left[\sum_{s=t+1}^{N} u(q_s, \alpha_s) - \tau^0(q_s, k_s) + \rho\left[\Delta\Pi + \sum_{s=t+1}^{N} \pi^0(q_s, k_s) - \overline{\pi}^0(k_s)\right]^+ \Big| h_t\right]$$

It is immediate that V is increasing in $\Delta\Pi$. In addition, at any history h_t, with current state and preferences α_t, k_t the agent's optimal consumption q_t solves

$$\max_{q_t \in Q} u(q_t, \alpha_t) - \tau^0(q_t, k_t) + V\left(\sum_{s=1}^{t} \pi^0(q_s, k_s) - \overline{\pi}^0(k_s)\right).$$

Hence it follows that the optimal policy at t is a function of α_t, k_t and cumulated excess profits $\Delta\Pi^t \equiv \sum_{s=1}^{t-1} \pi^0(q_s, k_s) - \overline{\pi}^0(k_s)$.

Given $\widehat{\sigma}$, consider any history h_t with associated excess profits $\Delta\Pi^t$. The next step of the proof establishes that if prescribed consumption $\widehat{q}$ under $\widehat{\sigma}$, differs from prescribed consumption q_0 under σ_0 at that history, it must be that $\pi^0(\widehat{q}, k_t) > \pi^0(q_0, k_t)$. Indeed, by optimality of $\widehat{q}$ under contract $\widehat{\tau}$, it follows that

$$\begin{aligned} u(\widehat{q}, \alpha_t) - \tau^0(\widehat{q}, k_t) &+ V\left(\Delta\Pi_t + \pi^0(\widehat{q}, k_t) - \overline{\pi}^0(k_t), h_t\right) \\ &> u(q_0, \alpha_t) - \tau^0(q_0, k_t) + V\left(\Delta\Pi_t + \pi^0(q_0, k_t) - \overline{\pi}^0(k_t), h_t\right). \end{aligned}$$

Since by definition of σ_0,

$$u(q_0, \alpha_t) - \tau^0(q_0, k_t) \geq u(\widehat{q}, \alpha_t) - \tau^0(\widehat{q}, k_t)$$

it follows that

$$V(\Delta\Pi_t + \pi^0(\widehat{q}, k_t) - \overline{\pi}^0(k_t), h_t) > V(\Delta\Pi_t + \pi^0(q_0, k_t) - \overline{\pi}^0(k_t), h_t).$$

By monotonicity of V in its first argument, this implies that

$$\pi^0(\widehat{q}, k_t) > \pi^0(q_0, k_t). \tag{A.7}$$

If prescribed consumption under $\widehat{\sigma}$ and σ_0 are the same, it is immediate that $\pi^0(\widehat{q}, k_t) = \pi^0(q_0, k_t)$.

Consider aggregate payoffs to the consumer. Recall that q_t^0 and $\widehat{q}_t$ denote consumption choices made under σ_0 and $\widehat{\sigma}$. Let η denote the consumer's payoff improvement from using contract $\widehat{\tau}$ instead of τ^0. By definition, using v^0 to denote net flow payoffs to the consumer under contract τ^0,

$$\mathbb{E}\left(\sum_{t=1}^{N} v^0(\widehat{q}_t, k_t, \alpha_t) + \rho\left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+\right) = \mathbb{E}\left(\sum_{t=1}^{N} v^0(q_t^0, k_t, \alpha_t)\right) + \eta$$

By optimality of σ_0 over $\widehat{\sigma}$ under τ^0,

$$\mathbb{E}\left(\sum_{t=1}^{N} v^0(q_t^0, k_t, \alpha_t)\right) \geq \mathbb{E}\left(\sum_{t=1}^{N} v^0(\widehat{q}_t, k_t, \alpha_t)\right).$$

Altogether, this implies that

$$\mathbb{E}\left(\left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+\right) \geq \frac{1}{\rho}\eta. \tag{A.8}$$

Consider now aggregate profits. Profits under σ_0, τ^0 and $\widehat{\sigma}$, $\widehat{\tau}$ respectively take the form

$$\Pi_N^0 = \sum_{t=1}^{N} \pi^0(q_t^0, k_t) \quad \text{and} \quad \widehat{\Pi}_N = \sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \rho\left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+.$$

It follows that

$$\begin{aligned}
\widehat{\Pi}_N - \Pi_N^0 &= \sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \pi^0(q_t^0, k_t) - \rho\left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+ \\
&= \sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t) + \sum_{t=1}^{N} \overline{\pi}^0(k_t) - \pi^0(q_t^0, k_t) - \rho\left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+ \\
&= \left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+ - \left[\sum_{t=1}^{N} \overline{\pi}^0(k_t) - \pi^0(\widehat{q}_t, k_t)\right]^+ + \sum_{t=1}^{N} \overline{\pi}^0(k_t) - \pi^0(q_t^0, k_t) \\
&\quad - \rho\left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+
\end{aligned}$$

Using (A.7), it follows that

$$\begin{aligned}
\widehat{\Pi}_N - \Pi_N^0 &\geq (1-\rho)\left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+ - \left[\sum_{t=1}^{N} \overline{\pi}^0(k_t) - \pi^0(q_t^0, k_t)\right]^+ + \sum_{t=1}^{N} \overline{\pi}^0(k_t) - \pi^0(q_t^0, k_t) \\
&\geq (1-\rho)\left[\sum_{t=1}^{N} \pi^0(\widehat{q}_t, k_t) - \overline{\pi}^0(k_t)\right]^+ - \left[\sum_{t=1}^{N} \pi^0(q_t^0, k_t) - \overline{\pi}^0(k_t)\right]^+ \qquad \text{(A.9)}
\end{aligned}$$

This establishes the left-hand side equality in (9). By definition of $\overline{\pi}^0(k_t)$,

$$\left(\sum_{t=1}^{T} \pi^0(q_t^0, k_t) - \overline{\pi}^0(k_t)\right)_{T\in\{1,\ldots,N\}}$$

is a supermartingale with increments bounded by $2||\pi^0||_\infty$. It follows from the Azuma-Hoeffding inequality that there exists M such that

$$\begin{aligned}
\mathbb{E}\left(\left[\sum_{t=1}^{N} \pi^0(q_t^0, k_t) - \overline{\pi}^0(k_t)\right]^+\right) &= \int_0^{+\infty} \mathsf{prob}\left(\sum_{t=1}^{N} \pi^0(q_t^0, k_t) - \overline{\pi}^0(k_t) \geq x\right) \mathrm{d}x \\
&\leq M\sqrt{N}.
\end{aligned}$$

Together with (A.8) and (A.9), this implies that

$$\frac{1}{N}\mathbb{E}\left[\widehat{\Pi}_N - \Pi_N^0\right] \geq \frac{1-\rho}{\rho}\eta - M\frac{1}{\sqrt{N}}.$$

This concludes the proof. ■

References

Allcott, H. (2011): "Rethinking real-time electricity pricing," *Resource and energy economics*, 33, 820–842.

Aubin, C., D. Fougere, E. Husson, and M. Ivaldi (1995): "Real-time pricing of electricity for residential customers: Econometric analysis of an experiment," *Journal of Applied Econometrics*, 10, S171–S191.

Azevedo, E. M. and E. Budish (2019): "Strategy-proofness in the large," *The Review of Economic Studies*, 86, 81–116.

Blackwell, D. (1956): "An Analog of the Minimax Theorem for Vector Payoffs," *Pacific Journal of Mathematics*, 6, 1–8.

Borenstein, S. and S. Holland (2005): "On the Efficiency of Competitive Electricity Markets with Time-Invariant Retail Prices," *RAND Journal of Economics*, 469–493.

Braverman, M. and S. Chassang (2022): "Data-driven incentive alignment in capitation schemes," *Journal of Public Economics*, 207, 104584.

Brot-Goldberg, Z. C., A. Chandra, B. R. Handel, and J. T. Kolstad (2017): "What does a deductible do? The impact of cost-sharing on health care prices, quantities, and spending dynamics," *The Quarterly Journal of Economics*, 132, 1261–1318.

Cabot, C. and M. Villavicencio (2024): "Second-best electricity pricing in France: Effectiveness of existing rates in evolving power markets," *Energy Economics*, 136, 107673.

Carroll, G. (2015): "Robustness and linear contracts," *American Economic Review*, 105, 536–563.

——— (2017): "Robustness and separation in multidimensional screening," *Econometrica*, 85, 453–488.

Centers for Medicare & Medicaid Services (2014): "Data Entrepreneurs' Synthetic Public Use File (DE-SynPUF) User Guide," Tech. rep., Centers for Medicare & Medicaid Services, synthetic Medicare claims data for research and software development.

Chassang, S. (2013): "Calibrated incentive contracts," *Econometrica*, 81, 1935–1971.

Chassang, S. and S. Kapon (2022): "Prior-free dynamic allocation under limited liability," *Theoretical Economics*, 17, 1109–1143.

de Clippel, G., P. Moscariello, P. Ortoleva, and K. Rozen (2024): "Caution in the face of complexity," Tech. rep., Mimeo.

Dizon-Ross, R. and A. D. Zucker (2025): "Mechanism design for personalized policy: A field experiment incentivizing exercise," Tech. rep., National Bureau of Economic Research.

Dworczak, P., S. D. Kominers, and M. Akbarpour (2021): "Redistribution through markets," *Econometrica*, 89, 1665–1698.

EDF (2025): "Tempo : Économies sur Facture Électricité," https://particulier.edf.fr/fr/accueil/gestion-contrat/options/tempo/details.html, accessed December 4^{th} 2025.

Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, and the Oregon Health Study Group (2012):

"The Oregon health insurance experiment: evidence from the first year," *The Quarterly journal of economics*, 127, 1057–1106.

Fowlie, M., C. Wolfram, P. Baylis, C. A. Spurlock, A. Todd-Blick, and P. Cappers (2021): "Default effects and follow-on behaviour: Evidence from an electricity pricing program," *The Review of Economic Studies*, 88, 2886–2934.

Greenberg, K., P. A. Pathak, and T. Sönmez (2024): "Redesigning the US Army's branching process: A case study in minimalist market design," *American Economic Review*, 114, 1070–1106.

Hannan, J. (1957): "Approximation to Bayes Risk in Repeated Play," *Contributions to the Theory of Games*, 3, 97–139.

Ho, K. and R. S. Lee (2023): "Health insurance menu design for large employers," *The RAND Journal of Economics*, 54, 598–637.

Huber, P. J. (1964): "Robust Estimation of a Location Parameter," *The Annals of Mathematical Statistics*, 73–101.

Ida, T., T. Ishihara, K. Ito, D. Kido, T. Kitagawa, S. Sakaguchi, and S. Sasaki (2024): "Dynamic targeting: Experimental evidence from energy rebate programs," Tech. rep., National Bureau of Economic Research.

——— (2025): "Choosing who chooses: Selection-driven targeting in energy rebate programs," *Econometrica*.

Ito, K., T. Ida, and M. Tanaka (2023): "Selection on welfare gains: Experimental evidence from electricity plan choice," *American Economic Review*, 113, 2937–2973.

Jessoe, K. and D. Rapson (2014): "Knowledge is (less) power: Experimental evidence from residential energy use," *American Economic Review*, 104, 1417–1438.

Joskow, P. L. and C. D. Wolfram (2012): "Dynamic pricing of electricity," *American Economic Review*, 102, 381–385.

Laffont, J.-J. and J. Tirole (1986): "Using cost observation to regulate firms," *Journal of political Economy*, 94, 614–641.

Lite (2023): "Tempo d'EDF : des tarifs inédits qui incitent fortement à la flexibilité," https://www.lite.eco/blog/tempo-edf-des-tarifs-inedits-qui-incitent-fortement-a-la-flexibilite/, accessed: 2025-11-26.

Madarász, K. and A. Prat (2017): "Sellers with misspecified models," *The Review of Economic Studies*, 84, 790–815.

Maskin, E. and J. Tirole (1990): "The principal-agent relationship with an informed principal: The case of private values," *Econometrica: Journal of the Econometric Society*, 379–409.

——— (1992): "The principal-agent relationship with an informed principal, II: Common values," *Econometrica: Journal of the Econometric Society*, 1–42.

Myerson, R. B. (1983): "Mechanism design by an informed principal," *Econometrica: Journal of the Econometric Society*, 1767–1797.

Nabil, T., G. Agoua, P. Cauchois, A. de Moliner, and B. Grossin (2025): "A Synthetic Dataset of French Electric Load Curves with Temperature Conditioning," in *Proceedings of the Workshop on Tackling Climate Change with Machine Learning at ICLR 2025*.

Newhouse, J. P. and Insurance Experiment Group (1993): *Free for All? Lessons From the Rand Health Insurance Experiment*.

Reguant, M. and M. Wagner (2025): "Smart Power Limits: Designing Shortage Mechanisms for Extreme Events," Tech. rep., National Bureau of Economic Research.

Reiss, P. C. and M. W. White (2005): "Household electricity demand, revisited," *The Review of Economic Studies*, 72, 853–883.

Rogerson, W. P. (2003): "Simple menus of contracts in cost-based procurement and regulation," *American Economic Review*, 93, 919–926.

RTE (2025): "Télécharger les indicateurs — éCO2mix données temps réel," https://www.rte-france.com/donnees-publications/eco2mix-donnees-temps-reel/telecharger-indicateurs, accessed: 2025-12-04.

Sönmez, T. (2023): "Minimalist market design: A framework for economists with policy aspirations," *arXiv preprint arXiv:2401.00307*.

Wolak, F. A. (2011): "Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment," *American Economic Review*, 101, 83–87.

——— (2019): "The role of efficient pricing in enabling a low-carbon electricity sector," *Economics of Energy & Environmental Policy*, 8, 29–52.