**FUNDAÇÃO GETULIO VARGAS**

**ESCOLA BRASILEIRA DE ECONOMIA E FINANÇAS**

**GABRIEL GRAL**

# REPEATED TRADE WITH DIVISIBLE GOODS

**RIO DE JANEIRO**

**2026**

**GABRIEL GRAL**

# REPEATED TRADE WITH DIVISIBLE GOODS

Dissertação apresentada à Escola Brasileira de Economia e Finanças da Fundação Getulio Vargas, como requisito para obtenção do título de Mestre em Economia.

Área de concentração: Teoria Econômica.

Orientador: Humberto Moreira.

**RIO DE JANEIRO**

**2026**

Dados Internacionais de Catalogação na Publicação (CIP)

Ficha catalográfica elaborada pelo Sistema de Bibliotecas/FGV

Gral, Gabriel Milliati

Repeated trade with divisible goods / Gabriel Milliati Gral. – 2026.

45 f.

Dissertação (mestrado) – Escola Brasileira de Economia e Finanças.

Orientador: Humberto Moreira

Inclui bibliografia.

1. Informação assimétrica. 2. Contratos. 3. Teoria dos jogos. 4. Seleção adversa (Seguros). I. Moreira, Humberto Ataíde II. Fundação Getulio Vargas. Escola Brasileira de Economia e Finanças. III. Título.

CDD – 330.01

Elaborada por Marcelle Costal de Castro dos Santos – CRB-7-RJ-007517/O

**FUNDAÇÃO GETULIO VARGAS**
**EPGE ESCOLA BRASILEIRA DE ECONOMIA E FINANÇAS - FGV EPGE**
**MESTRADO EM ECONOMIA**

# FOLHA DE APROVAÇÃO

**GABRIEL MILLIATI GRAL**

**"REPEATED TRADE WITH DIVISIBLE GOODS".**

**DISSERTAÇÃO APRESENTADO/A AO CURSO DE MESTRADO EM ECONOMIA QUE CONFERIU O GRAU DE MESTRE EM ECONOMIA.**

**DATA DA DEFESA: 15/01/2026**

**ASSINATURA DOS MEMBROS DA BANCA EXAMINADORA**

**PRESIDENTE DA COMISSÃO EXAMINADORA: PROF. HUMBERTO LUIZ ATAÍDE MOREIRA**

*<ASSINADO ELETRONICAMENTE>*
**PROF. HUMBERTO LUIZ ATAÍDE MOREIRA**
**ORIENTADOR**

*<ASSINADO ELETRONICAMENTE>*
**PROF. WILLIAM FUCHS**
**MEMBRO EXT - UT AUSTIN**

*<ASSINADO ELETRONICAMENTE>*
**PROF. LUCAS JÓVER MAESTRI**
**MEMBRO INT. - FGV EPGE**

RIO DE JANEIRO, 15 DE JANEIRO DE 2026.

*<ASSINADO ELETRONICAMENTE>*
**PROF. LUCAS JÓVER MAESTRI**
**COORDENADOR**

*<ASSINADO ELETRONICAMENTE>*
**PROF. ANTONIO DE ARAUJO FREITAS JUNIOR**
**PRÓ-REITOR DE ENSINO, PESQUISA E PÓS-GRADUAÇÃO**

ESTE É UM TRABALHO ORIGINAL ONDE FOI VERIFICADA A NÃO EXISTÊNCIA DE PLÁGIO E DE UTILIZAÇÃO DE INTELIGÊNCIA ARTIFICIAL, NÃO EXPLICITADA, NO CORPO DO TRABALHO ATESTADO PELO ALUNO(A) E ORIENTADOR(A). ESTE DOCUMENTO NÃO CONFERE TÍTULO. PARA TAL DEVERÃO SER CUMPRIDOS OS REQUISITOS DO PROGRAMA DE PÓS-GRADUAÇÃO.

**D4Sign** 7c447066-eeae-4814-aca7-03f32e9c5ac3 - Para confirmar as assinaturas acesse https://secure.d4sign.com.br/verificar
**Documento assinado eletronicamente, conforme MP 2.200-2/01, Art. 10º, §2. Brasil**

*À Santa Mãe de Deus*

# Agradecimentos

Agradeço ao Humberto Moreira pela orientação excepcional, pela parceria e pela confiança ao longo deste trabalho. Também tiveram papel fundamental em minha formação e no desenvolvimento desta dissertação os professores William Fuchs, Lucas Maestri e Aloísio Araújo.

Meus pais, Anselmo e Márcia, e meu irmão caçula e economista predileto, Rafael, foram meu apoio e meu motivo para cada conquista até aqui. As amizades que construí nos últimos anos tornaram o caminho divertidíssimo, e o Rio de Janeiro deixou tudo ainda mais bonito.

Esta dissertação foi financiada, em parte, pela Fundação Getulio Vargas (FGV), pelo Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) e pela Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro (FAPERJ).

## RESUMO

Analisamos um problema de seleção adversa em negociação repetida com contratos divisíveis e valores interdependentes. Um indivíduo informado encontra sequencialmente contrapartes de vida curta. O indivíduo escolhe uma quantidade para negociar, e as contrapartes competem em preços após observar todo o histórico de negociações. Propomos uma regra simples de formação de crenças e mostramos que, sob separação, ela é equivalente ao refinamento D1. O equilíbrio é único, separador e dependente da história. Nos períodos iniciais, os incentivos dinâmicos ampliam a distorção estática: antecipando que tipos piores podem imitá-los para obter preços melhores no futuro, quase todos os tipos de indivíduos distorcem mais as quantidades em relação ao *benchmark* estático. À medida que o histórico público se acumula, o valor marginal de sinalização adicional diminui, e as alocações se movem monotonamente em direção ao *first-best*. Conforme o horizonte do jogo aumenta, as alocações dos períodos finais convergem para a alocação eficiente. Mostramos que a repetição das negociações enfraquece endogenamente o papel sinalizador da retenção ao longo do tempo. Ilustramos os resultados em três aplicações: seguros, financiamento e mercado de trabalho.

**Palavras-chave:** seleção adversa; teoria dos contratos; jogos dinâmicos.

## ABSTRACT

We study adverse selection in repeated trade with divisible contracts and interdependent values. A privately informed individual sequentially meets short-lived dealers, posts a quantity, and dealers compete on prices after observing the full trading history. We propose a simple belief formation rule and show that, under separation, it is equivalent to D1 in this dynamic setting. The equilibrium is unique, separating and history-dependent. Early on, dynamic incentives amplify the static distortion: anticipating that worse types might mimic them to secure better prices in the future, almost all types scale back quantities relative to the static benchmark. As public history accumulates, the marginal value of additional signaling falls, and allocations move monotonically toward the first-best. As the game horizon lengthens, terminal allocations converge to efficiency. Our work clarifies how quantities and timing interact as signals, showing that repeated trade endogenously weakens the signaling role of retention over time. We illustrate the results in three applications: insurance, financing and the job market.



**Keywords:** adverse selection; contract theory; dynamic games.

# Contents


**1 Introduction** **1**

**2 Related Literature** **3**

**3 Model** **5**

**4 Applications** **7**
- 4.1 Insurance . . . 7
- 4.2 Financing . . . 8
- 4.3 Job Market . . . 8

**5 Benchmark: Commitment to Long-Term Contracts** **10**

**6 Repeated Trade Without Commitment** **14**
- 6.1 The $\Gamma$ Criterion . . . 17
- 6.2 The D1 Equilibrium . . . 19
- 6.3 Observability . . . 22

**7 Applications: Results** **24**
- 7.1 Insurance . . . 24
- 7.2 Financing . . . 24
- 7.3 Job Market . . . 25

**8 Extension: Capacity Constraint** **27**
- 8.1 Adapted Model . . . 27
- 8.2 Commitment . . . 27
- 8.3 No Commitment . . . 28

**9 Conclusion** **30**

**References** **31**

**Appendix: Proofs** **35**

# 1 Introduction

Retention and timing are classic signaling devices in markets with adverse selection. A privately informed individual can either distort quantities or delay contracting to convey quality.[1] We study a setting with interdependent values and a continuum of types where both levers interact: contracting repeats over a finite horizon, and the individual chooses a quantity each period.

To ground the discussion, consider insurance as a running example. Each year an individual privately informed about risk chooses a level of coverage, and short-lived insurers, observing that choice, compete on price. In the static game, safer types buy less coverage, reveal themselves, and obtain fair premiums. With repeated spot contracting, however, simply replaying the static allocation is not incentive compatible. A clever high-risk type would cut coverage further today to appear safer tomorrow and obtain lower future premiums. This intertemporal signaling motive makes the analysis nontrivial: insurers' beliefs following any history must account for a type-dependent continuation strategy.

We propose a simple identification discipline, which we call the $\Gamma$-rule, that maps any observed trade history to types via a one-dimensional statistic of past and current quantities. Under separation, beliefs induced by this rule coincide with those that satisfy the D1 refinement of Cho and Kreps (1987); in particular, any nonzero trade history is assigned to a single type. This equivalence yields a tractable characterization of equilibrium as a system of coupled first-order differential equations, with boundary conditions pinned down by an extreme type. The analysis clarifies how dynamic incentives shape quantities and prices along the path, and how retention and timing interact as signals in repeated markets.

Equilibrium unfolds as follows. If the individual could commit at the outset to the full sequence of trades, the time structure then plays no signaling role: quantities are time-invariant and replicate the static allocation each period, providing a natural benchmark. With spot contracting, dynamics bite. The equilibrium is still unique and separating, but history-dependent. Early on, almost all types scale back quantities relative to the static benchmark to deter mimicry by worse types. As public history accumulates, the marginal value of further signaling falls, distortions shrink, and allocations move monotonically toward the stage first-best. As the game horizon lengthens, allocations of terminal periods converge to efficiency.

Transparency of trades is pivotal. If past trades are opaque, each period reduces to a one-shot interaction and the commitment allocation is implemented. With observable histories, early signaling is costly but later allocations improve. Welfare comparisons are thus path-dependent: opacity is preferred ex ante, whereas ongoing transparency is desirable once its initial costs are sunk.

[1] Among many others, Leland and Pyle (1977) explore retention and Fuchs and Skrzypacz (2019) explore timing as signals.

Our results apply broadly to repeated-trading environments with adverse selection. We highlight three applications. First, in our motivating insurance setting, cumulative coverage serves as the sufficient statistic to convey quality. Low-risk types initially retain more risk to credibly separate and then converge monotonically to full insurance. Second, in financing decisions with risk-averse dealers, a firm repeatedly chooses how much of an asset to sell. Types are identified by total discounted sales, so higher-quality assets are sold less initially and quantities rise over time. Third, in a stylized job market, a worker privately informed about output quality offers hours each period. Depending on how quality correlates with effort costs, higher-quality workers initially distort hours relative to the symmetric information case, and then adjust toward efficiency over time.

As an extension, we study intertemporal capacity constraints. This occurs, for instance, when a unit asset is sold in parts over time. In this case, past trades serve not only to signal quality but also to inform dealers about the seller's residual supply, which constrains future feasible offers. Residual supply interacts with dynamic incentives and complicates off-path beliefs. Nonetheless, we can characterize the equilibrium for a game with two periods. We find a two-phase equilibrium: one interval of types operates under binding capacity, while the complementary interval is slack and behaves as under the $\Gamma$-rule. This links our analysis to Fuchs, Gottardi, and Moreira (2025) and shows how curvature can alter their conclusions.

The paper proceeds as follows. Section 2 situates our contribution in the literature. Section 3 presents the model. Section 4 introduces the applications and relates each to its literature. Section 5 analyzes the commitment benchmark. Section 6 studies the no-commitment case, introduces our identification discipline, proves its equivalence to D1, and characterizes the equilibrium paths. Section 7 interprets the results within each application. Section 8 extends the analysis to capacity constraints. Section 9 offers concluding remarks. Formal proofs for all lemmas and propositions can be found in the Appendix.

# 2 Related Literature

We connect two strands of work: *(i)* dynamic adverse selection, which emphasizes intertemporal incentives and learning from publicly observed histories; and *(ii)* adverse selection with divisible trades, where quantity choices convey private information and affect pricing.

First, we relate to dynamic extensions of Akerlof (1970). Noldeke and Van Damme (1990) and Swinkels (1999) are early dynamic extensions of the signaling solution of Rothschild and Stiglitz (1976). Vincent (1990) studies dynamic auctions in which two uninformed buyers compete for a durable good held by an informed seller. These papers find that higher-quality trades occur later in equilibrium. We share a dynamic and competitive environment, but introduce *divisible* contracting. We also build on a modern literature on dynamic signaling that relies on parsimonious models of when to post an asset for sale, including Janssen and Roy (2002), Kremer and Skrzypacz (2007), and Fuchs, Öry, and Skrzypacz (2016). None of these works features divisibility. Complementary lines study dynamics driven by information arrival (Daley and Green, 2012; Kaya and Kim, 2018), as well as optimal intervention in dynamic selection markets (Fuchs and Skrzypacz, 2015). More recently, Du and Zhu (2017) and Fuchs and Skrzypacz (2019) study how alternative timing structures shape trade in lemons markets.

Second, on divisibility, the insight of Leland and Pyle (1977) is central: issuers may retain part of an asset to signal quality and secure better prices. Related results appear in the security-design literature (DeMarzo and Duffie, 1999). In our setting, types similarly distort quantities each period to separate. With repeated trade, we show that distortions are intensified early on and attenuate over time.

Closest to our work are papers on interactions between a long-lived informed agent and a sequence of short-lived uninformed dealers. Like us, Fuchs, Gottardi, and Moreira (2025) study the interaction between retention and timing as signals, but they impose linear preferences and an intertemporal resource constraint. Their setting is therefore one in which a seller allocates a single divisible good over time, whereas ours features repeated trading of a perishable good in each period. They show that, under long-term contracts, only discounted quantity matters as a signal. Under spot contracting, timing itself becomes a commitment device: by delaying trade and then selling the entire residual supply, the informed party effectively mimics commitment, so the no-commitment outcome can replicate the commitment allocation.[1] Kaya and Roy (2024) also study repeated trading and the role of accumulated information for efficiency under monopoly and competition.[2] Our model differs by allowing dealers to observe each seller's full history rather than only aggregate volume, by accommodating a continuum of types, and, crucially, by permitting divisible trades.

[1] In our capacity-constraint extension (Section 8), concavity in quantities breaks this equivalence: the commitment allocation need not be replicated without commitment.

[2] Related work on transparency includes Dilmé (2025), who studies disclosure design in auctions where buyers' willingness to pay is independent of seller type.

Our discussion of commitment relates to prior work on repeated or dynamic negotiations with and without commitment. Such issues arise naturally in the insurance literature, which includes Rubinstein and Yaari (1983), Hosios and Peters (1989), Rey and Salanie (1990), Hendel and Lizzeri (2003), and Daily, Hendel, and Lizzeri (2008). The absence of commitment can also be interpreted as non-exclusive contracting, thus relating to Attar, Mariotti, and Salanié (2011) and Asriyan and Vanasco (2024). Our focus, however, is on non-exclusivity across dates – via lack of commitment and potentially opaque histories – rather than on non-exclusivity within a single trading date.

More broadly, our analysis connects to dynamic adverse selection in other market structures. With two long-lived parties, Hart and Tirole (1988) highlight ratchet-effect forces that can stall learning absent long-term commitment. In bargaining with a divisible good and interdependent values, Gerardi, Maestri, and Monzón (2022) find trade patterns in which high-types sell smoothly over time while low-types may liquidate at once, resonating with our finding that better types sustain larger, longer-lived distortions. A wider bargaining literature similarly documents that higher-quality goods are traded more slowly (Evans, 1989; Deneckere and Liang, 2006; Hörner and Vieille, 2009). Additional connections include dynamic adverse selection with search (Guerrieri, Shimer, and Wright, 2010; Camargo and Lester, 2014) and a growing literature on dynamic mechanism design (Courty and Hao, 2000; Eső and Szentes, 2007; Pavan, Segal, and Toikka, 2014).

This review is not exhaustive, and we provide further references when discussing each application in Sections 4 and 7.

# 3 Model

We model the dynamic interaction of a long-lived individual with a competitive market of short-lived dealers. The individual privately knows his type $\theta \in [0,1]$, drawn from a publicly known distribution with full support. Trading occurs on a finite number of evenly spaced dates $t \in \{0,1,\ldots,T\}$.

At each date $t$, the individual announces a quantity $x_t \in \mathbb{R}_+$, representing the amount he wishes to trade in that period. Dealers indexed by $i \in \mathcal{I}_t$, each short-lived and active only at $t$, compete by submitting price quotes $p_{it}$ in response to $x_t$, with $2 \leq |\mathcal{I}_t| \leq \infty$. The individual observes all quotes and either accepts the best one, forming the contract $(p_t, x_t)$, or rejects all and proceeds to the next date without contracting.

The individual's instantaneous utility from $(p_t, x_t)$ is

$$w(\theta, p_t, x_t) \coloneqq f(p_t, x_t) + \theta\, g(x_t), \tag{3.1}$$

with $f$ and $g$ twice differentiable. We normalize the outside option to zero for all types, so $w$ is net utility relative to the outside option. Thus, rejecting all offers at $t$ corresponds to the null contract $(0,0)$. The individual discounts future utility at rate $\beta \in (0,1)$, so his total utility from a sequence $\{(p_t, x_t)\}_{t=0}^{T}$ is $\sum_{t=0}^{T} \beta^t w(\theta, p_t, x_t)$.

Dealer $i$'s expected payoff from $(p_{it}, x_t)$ is $\mathbb{E}\left[\pi(\theta, p_{it}, x_t)\right]$, where $\pi$ is twice differentiable and the expectation is taken with respect to her belief about $\theta$, conditional on the information available at time $t$. To capture a conflict of interest over price, we assume $w_2(\theta, p_t, x_t)$ and $\pi_2(\theta, p_t, x_t)$ have opposite, nonzero signs on the relevant domain.[1] The best price following an offer $x_t$ is $p_t = \min_{i \in \mathcal{I}_t} p_{it}$ when $w_2 < 0$, and $p_t = \max_{i \in \mathcal{I}_t} p_{it}$ when $w_2 > 0$.

Under symmetric information, perfect competition implies a zero-profit condition that implicitly defines a pricing function $\mathsf{p}(\theta, x_t)$ by $\pi(\theta, \mathsf{p}(\theta, x_t), x_t) = 0$. By the price-conflict and the smoothness of $\pi$, the implicit function theorem ensures $\mathsf{p}$ is well-defined and differentiable in a neighborhood of interest. Back to the asymmetric information case, we can use $\mathsf{p}$ to internalize pricing in the utility function via beliefs:

$$u(\theta, \hat{\theta}, x_t) \coloneqq w(\theta, \mathsf{p}(\hat{\theta}, x_t), x_t).$$

The function $u$ inherits differentiability from $w$ and $\pi$. We use $u$ and $w$ as needed, depending on whether we emphasize how choices affect beliefs (and hence prices) or prices directly.

ASSUMPTION 1. The function $u$ satisfies:

1.A. $u_{1,3}(\theta, \hat{\theta}, x_t) = g'(x_t)$ never equals zero, so it has constant sign; *(type monotonicity)*

1.B. $u_2(\theta, \hat{\theta}, x_t)$ never equals zero, so it has constant sign; *(belief monotonicity)*

[1] When no ambiguity arises, subscripts denote partial derivatives.

1.C. $u_{3,3}(\theta, \theta, x_t) < 0$ for all $x_t$, and the problem $\max_{x_t} u(\theta, \theta, x_t)$ has a positive finite solution $x^{FB}(\theta)$ for all $\theta \in [0, 1]$, satisfying $u_3\left(\theta, \theta, x^{FB}(\theta)\right) = 0$.

The signs of type and belief monotonicity vary across applications, and we preview their roles. Type monotonicity coincides with equilibrium allocation slopes: $x_t(\theta)$ increases with $\theta$ in each period when type monotonicity is positive ($u_{1,3} = g'(x_t) > 0$), and decreases otherwise. Belief monotonicity pins down the boundary condition: the boundary corresponds to efficient contracting for type $0$ when belief monotonicity is positive ($u_2 > 0$), and for type $1$ otherwise. We often refer to this type as the *worst type* (i.e., type $0$ when $u_2 > 0$ and type $1$ otherwise). We refer to $x^{FB}(\theta)$ as the *stage first-best* quantity. When the horizon extends beyond a single period ($T > 0$), the first-best allocation for each type repeats the stage first-best at every date.

To ensure existence and uniqueness of equilibrium, we impose:

ASSUMPTION 2. The function $u$ satisfies:

2.A. $|u_2(\theta, \theta, x_0)|$ is bounded;

2.B. $\frac{u_3(\theta,\theta,x_t)}{u_{1,3}(\theta,\theta,x_t)}$ is strictly increasing in $x_t$ when $u_{1,3} < 0$ and strictly decreasing otherwise;

2.C. $u_{1,3}(\theta, \theta, x_t)u_{2,3}(\theta, \theta, x_t) \leq 0$ for all $x_t$.

Assumption 2.B places discipline on $g(x_t)$, imposing a stronger curvature condition. When $g(x_t)$ is linear, as in our applications, 2.B follows from 1.C. We state it to highlight that the framework accommodates more general forms. Assumption 2.C captures a conflict over quality: when marginal gains in quantity rise with the true type, they must fall with the perceived type, and vice versa.[2]

Throughout, write $x = (x_0, \ldots, x_T)$ for contracted quantities and $p = (p_0, \ldots, p_T)$ for prices. For any sequence $z$, let $z^t$ denote its truncation up to period $t$.

[2] We could relax assumption 2.C and still obtain uniqueness if the ratio $|u_{2,3}/u_{1,3}|$ were small enough, but we adopt simpler sufficient conditions.

# 4 Applications

We apply our framework to three settings. The first, insurance, is our running example and anchors the intuition in the equilibrium analysis that follows. For clarity, each application adopts simple functional forms for $w$ and $\pi$, which we also use in the numerical illustrations. Alternative primitives are admissible provided they respect the structure in Section 3. We return to these applications in Section 7 to interpret the general results, present the numerics, and connect the results to the relevant literature.

## 4.1 Insurance

Insurance contracts are typically written for a finite term, so renegotiation naturally recurs. Canonical models (Rothschild and Stiglitz, 1976; Wilson, 1977) have been extended to repeated contracting to examine renegotiation (Dionne, 1981; Dionne and Doherty, 1994), insurer learning under monopoly (Rubinstein and Yaari, 1983; Hosios and Peters, 1989), and commitment to long-term contracts (Hendel and Lizzeri, 2003). Our setting differs in two essential ways: learning is driven solely by publicly observable coverage choices and premiums (accident histories are unobserved), and the policyholder can choose any coverage level each period rather than a discrete contract.

We adopt a dynamic CARA-Gaussian setting.[1] An individual purchases insurance from risk-neutral, short-lived insurers (our "dealers"). His utility over wealth $y_t$ is $-\exp(-\gamma\, y_t)$ and wealth shocks follow $Z \sim \mathcal{N}(\theta, \sigma^2)$, where $\gamma > 0$ and $\sigma^2 > 0$ are publicly known. A contract $(p_t, x_t)$ specifies premium $p_t$ and coverage $x_t$.[2] Normalizing initial wealth to zero, final wealth at $t$ is $y_t = -p_t - (1 - x_t)Z$, so expected utility is

$$\mathbb{E}\left[-\exp(-\gamma\, y)\right] = -\exp\left\{-\gamma\left(-p_t - (1-x_t)\,\theta\right) + \tfrac{\gamma^2}{2}(1-x_t)^2\sigma^2\right\}.$$

Maximizing expected utility is equivalent to maximizing the certainty equivalent

$$\text{CE}(\theta, p_t, x_t) = -p_t - (1-x_t)\theta - \tfrac{\gamma\sigma^2}{2}(1-x_t)^2.$$

Subtracting $\text{CE}(\theta, 0, 0)$ yields

$$w(\theta, p_t, x_t) = -p_t - \tfrac{\gamma\sigma^2}{2}(x_t^2 - 2x_t) + \theta\, x_t,$$

which matches (3.1). Insurers' profits are $\pi(\theta, p_{it}, x_t) = p_{it} - \theta\, x_t$, implying the zero-profit pricing function $\mathsf{p}(\theta, x_t) = x_t\,\theta$. In $u$-notation,

$$u(\theta, \hat{\theta}, x_t) = x_t\,(\theta - \hat{\theta}) - \tfrac{\gamma\sigma^2}{2}(x_t^2 - 2x_t).$$

[1] The closest specification is Levy and Veiga (2025). Related settings include Einav et al. (2013); Azevedo and Gottlieb (2017) and Farinha Luz, Gottardi, and Moreira (2023).

[2] Results are unchanged if coverage is restricted to $x_t \in [0, 1]$; we keep $x_t \in \mathbb{R}_+$ for consistency with the general setup.

Type monotonicity is positive: $u_{1,3} = 1 > 0$. Belief monotonicity is negative: $u_2 = -x_t < 0$. The stage first-best features full-insurance for all types: $x^{FB}(\theta) = 1$.

## 4.2 Financing

Beyond Leland and Pyle (1977), classic finance papers study retention as a signal. Myers and Majluf (1984) show that privately informed managers may pass up valuable projects rather than issue information-sensitive equity. DeMarzo and Duffie (1999) embed retention in security design and show that, under certain conditions, standard debt is optimal because it is least information-sensitive. More recently, Li (2022) considers unverifiable retention, which relates to our notion that the informed party may not retain as dealers expect in future periods. Our application parallels these ideas in a repeated trading environment. The financing decision – selling a fraction of the asset – is simple, but repetition generates rich dynamic signaling and pricing.

A firm sells a divisible asset and privately knows its quality $\theta \in [0, 1]$. The period-$t$ payoff from selling quantity $x_t$ at price $p_t$ is $w(\theta, p_t, x_t) = p_t - \theta x_t$, so the structure in (3.1) applies. Interpret $\theta$ as either average production cost or the owner's private benefit from retention. Curvature arises from dealers' risk aversion (rather than the individual's, as in insurance). In a CARA-Gaussian setup, dealer $i$ has utility $-\exp(-\gamma y_t)$ over wealth $y_t$.[3] If she pays $p_{it}$ for $x_t$, her wealth is $y_t = x_t(v(\theta) + Z) - p_{it}$. The valuation function $v$ is continuously differentiable with $v'(\theta) > 0$ and $v(\theta) > \theta$ for all $\theta \in [0, 1)$, valuation shocks follow $Z \sim \mathcal{N}(0, \sigma^2)$, and $\gamma > 0$ and $\sigma^2 > 0$ are publicly known.

Under full information, dealer $i$'s certainty equivalent for contract $(x_t, p_{it})$ is

$$\text{CE}(\theta, p_t, x_t) = x_t v(\theta) - \tfrac{\gamma\sigma^2}{2} x_t^2 - p_{it}.$$

The zero-profit pricing function is $\mathsf{p}(\theta, x_t) = x_t v(\theta) - \frac{\gamma\sigma^2}{2} x_t^2$, which yields the reduced-form payoff

$$u(\theta, \hat{\theta}, x_t) = x_t(v(\hat{\theta}) - \theta) - \tfrac{\gamma\sigma^2}{2} x_t^2.$$

Type monotonicity is negative (since here $g(x_t) = -x_t$), and belief monotonicity is positive.

## 4.3 Job Market

The static signaling model of Spence (1973) for the job market was extended to dynamic offers by Noldeke and Van Damme (1990), who consider public offers, and Swinkels (1999), who consider private offers. Hörner and Vieille (2009) further compare private versus public bargaining, and Kremer and Skrzypacz (2007) study dynamics with public revelation of information. This application shares the dynamic feature, but differs by allowing contracts to be divisible in hours. Thus the labor contract itself, in the form of hours and wages, acts as the signal rather than education or grades. Equivalently, our analysis can be read as conditioning on education and adding the history of contracts as an extra signaling margin.

[3] See Biais, Martimort, and Rochet (2000) for a similar setting.

Empirically, Blair et al. (2020) document that the content of current work can function as a job market signal, especially for non-college-educated workers. Here, the high-dimensional content is distilled to the history of work contracts.

Each period the worker offers $x_t$ hours. His utility from salary $p_t$ is $w(\theta, p_t, x_t) = p_t - c(\theta, x_t)$, and firms earn $\pi(\theta, p_t, x_t) = v(\theta)x_t - p_t$. Hence the reduced-form utility is

$$u(\theta, \hat{\theta}, x_t) = v(\hat{\theta})x_t - c(\theta, x_t).$$

Assume $v$ satisfies the same conditions as in the previous application, and $c$ is smooth, increasing, and convex in $x_t$. Belief monotonicity is positive: $u_2 = x_t v'(\hat{\theta}) > 0$. Type monotonicity depends on the cross-partial of costs: $u_{1,3} = -c_{1,2}(\theta, x_t)$. In accordance with Assumption 2, we assume better types dislike hours more. A possible specification is $c(\theta, x_t) = \theta x_t + x_t^2$. Type monotonicity is thus negative.

# 5 Benchmark: Commitment to Long-Term Contracts

We begin our equilibrium analysis by considering a case where the individual can commit at the initial date to long-term contracts and cannot contract thereafter. This benchmark clarifies the roles of commitment and transparency. We first define an equilibrium, then prove four building blocks: zero-profits on path, separation, belief identification under D1, and time invariance of allocations. These are Lemmas 1 to 4, and they reduce the problem to a static, single-dimensional signaling game. Proposition 1 then characterizes the equilibrium path.

The individual commits at $t = 0$ to a path of offers $x = (x_0, \ldots, x_T)$, observed by dealers at all trading dates, so dealers may condition their quotes on the entire path. Let $\Omega := \mathbb{R}^{T+1}$ denote the space of long-term contracts. The pricing strategy of dealer $i \in \mathcal{I}_t$ is $p_{it}(x)$, and $p_t(x)$ is the best submitted quote at date $t$ given $x$. For offer path $x$ and induced prices $p(x)$, type $\theta$'s acceptance decision for the $(p_t(x), x_t)$ contract is $a_t(\theta, p(x), x) \in \{0, 1\}$, and his discounted utility is

$$W(\theta, x) = \sum_{t=0}^{T} \beta^t \, w(\theta, p_t(x), x_t) \, a_t(\theta, p(x), x).$$

This commitment environment parallels an Arrow-Debreu economy with date-specific goods: a unit at date $t$ is distinct from a unit at $t'$. Allowing arbitrary off-path beliefs generally yields multiple equilibria in signaling environments like ours. Perfect Bayesian equilibrium is too weak here because many off-path contracts are not Bayes-reachable, leaving posteriors largely unrestricted. To discipline beliefs we impose the D1 refinement of Cho and Kreps (1987), extended to a continuum of types by Ramey (1996). Intuitively, upon observing any off-path contract $x$, dealers concentrate beliefs on the types for whom deviating to $x$ is relatively most profitable. We restrict attention to pure strategies: each type $\theta$ chooses to offer $x^c(\theta)$ with certainty.

DEFINITION 1 (EQUILIBRIUM WITH COMMITMENT). A sequential equilibrium satisfying D1 in the market with commitment is a tuple $(x^c, \{\mu^c_{it}\}, \{p^c_{it}\}, \{a^c_t\})$ consisting of: a contract selection strategy $x^c(\theta)$; a system of dealer beliefs $\mu^c_{it}(x)$ assigning probability distributions over types conditional on the observed contract $x$; dealer pricing strategies $p^c_{it}(x)$; and individual acceptance strategies $a^c_t(p, x, \theta)$, such that:

(*1.1*) *Individuals optimize:* for each type $\theta$, contract selection and acceptance strategies satisfy

$$x^c(\theta) \in \arg\max_{x \in \Omega} W(x, \theta).$$

(*1.2*) *Dealers optimize:* taking the strategies of the individual and other dealers as given, each dealer $i$ chooses $p^c_{it}(x)$ to maximize her expected payoff at date $t$, conditional on $\mu^c_{it}(x)$.

(*1.3*) *On-path beliefs:* let $\Omega' \subseteq \Omega$ be the set of on-path contracts, $\Omega' := \{x^c(\theta) : \theta \in [0, 1]\}$. For any $x \in \Omega'$, dealers' beliefs are updated according to Bayes' rule whenever applicable. If

a contract is selected by a measure-zero set of types, belief support is only required to lie within that set.

(*1.4*) *Off-path beliefs (D1):* for any $x \in \Omega \setminus \Omega'$, a type $\theta$ is assigned zero posterior probability whenever, for any dealers' response $(p_0, \ldots, p_T)$, there exists another type $\theta'$ with a stronger incentive to deviate from equilibrium. That is, type $\theta'$ strictly prefers to deviate for any price response that gives type $\theta$ only a weak incentive to deviate.

This equilibrium concept builds the D1 restriction directly into the definition for off-path contracts while preserving standard Bayesian updating on path.[1] Given $x^c(\theta)$, let $K(x)$ denote the set of types that offer contract $x$ in the commitment equilibrium.

Under perfect competition in the dealers' market, dealers cannot earn strictly positive expected profits on-path by standard Bertrand logic.

LEMMA 1. *In any equilibrium* $(x^c, \{\mu_{it}^c\}, \{p_{it}^c\}, \{a_t^c\})$*, whenever* $K(x) \neq \emptyset$*, any accepted price* $p_t^c(x)$ *must yield zero expected profits for dealers at date* $t$*:*

$$\mathbb{E}\left[\pi(\theta, p_t^c(x), x_t) \,|\, \theta \in K(x)\right] = 0.$$

Now fix a commitment equilibrium $(x^c, \{\mu_{it}^c\}, \{p_{it}^c\}, \{a_t^c\})$ and let $U(\theta)$ denote the individual's equilibrium value function, defined as

$$U(\theta) \coloneqq \max_{x \in \Omega} W(x, \theta) = \sum_{t=0}^{T} \beta^t \left[f(p_t^c(x^c(\theta)), x_t^c(\theta)) + \theta g(x_t^c(\theta))\right]. \tag{5.1}$$

Define $\Gamma_T : [0,1] \to \mathbb{R}$ by

$$\Gamma_T(\theta) \coloneqq \sum_{t=0}^{T} \beta^t g\left(x_t^c(\theta)\right).$$

This function captures the component of the individual's on-path utility that depends directly on his type, which is clear in definition (5.1). By standard arguments, $U(\theta)$ is convex, and by the envelope condition, $U'(\theta) = \Gamma_T(\theta)$ almost everywhere. It follows that $\Gamma_T$ is non-decreasing in $\theta$. In the dynamic insurance setting, $\Gamma_T(\theta)$ equals the total discounted coverage purchased in equilibrium by type $\theta$, and monotonicity means that riskier types purchase greater coverage in equilibrium. The next results use $\Gamma_T$ extensively.

LEMMA 2. *Any commitment equilibrium that satisfies D1 must be separating.*

We provide intuition for the proof using the insurance case.[2] Suppose, towards a contradiction, that a set of types pools on one contract. By single-crossing, the pooled set is an interval. Competition implies on-path prices break even on that interval. Since the zero-profit price increases in type, the pooled prices exceed the break-even price for the lowest pooled type. Now

[1] Related belief discipline appears in competitive adverse-selection models such as Azevedo and Gottlieb (2017) and Bisin and Gottardi (2006).

[2] In the insurance setting we have $\partial w/\partial p < 0$, $\partial w/\partial \theta > 0$, $\partial \pi/\partial \theta < 0$.

take that lowest pooled type and marginally decrease one component of his contract. Whether the deviation is treated as on- or off-path, D1 shifts beliefs toward lower types, which delivers a strictly more favorable price vector than the pooling price. As the quantity change tends to zero while the discrete price improvement persists across dates, pooling is not sustainable. Hence any D1-refined commitment equilibrium must be separating.

In a separating equilibrium, the function $\Gamma_T$ must be strictly increasing. We have already established that it is weakly increasing. If it were constant over some interval, then all types in that interval would prefer to mimic the strategy of the lowest type in the interval, violating incentive compatibility. Hence $\Gamma_T$ is strictly monotone and admits a well-defined inverse function, which will be useful to characterize off-path beliefs.

LEMMA 3. *Consider any separating equilibrium* $(x^c, \{\mu_{it}^c\}, \{p_{it}^c\}, \{a_t^c\})$ *and its associated function* $\Gamma_T(\theta)$. *For any contract* $\bar{x}$, *define* $\Sigma_T(\bar{x}) := \sum_{t=0}^{T} \beta^t g(\bar{x}_t)$. *Dealers' beliefs upon observing* $\bar{x}$ *must concentrate on the singleton type*

$$\hat{\theta}(\bar{x}) = \begin{cases} \Gamma_T^{-1}\left(\Sigma_T(\bar{x})\right) & \textit{if } \Sigma_T(\bar{x}) \in [\Gamma_T(0), \Gamma_T(1)], \\ 0 & \textit{if } \Sigma_T(\bar{x}) < \Gamma_T(0), \\ 1 & \textit{if } \Sigma_T(\bar{x}) > \Gamma_T(1). \end{cases}$$

The on-path case is immediate. For an off-path offer, the claim is that beliefs concentrate on the unique $\hat{\theta}$ with $\Gamma_T(\hat{\theta}) = \Sigma_T(\bar{x})$. Our proof uses D1 and the convexity of the value function. Fix any other type $\theta' \neq \hat{\theta}$ and consider a price vector that makes $\theta'$ just indifferent between $\bar{x}$ and his equilibrium contract $x^c(\theta')$. By convexity, there exists a nearby type who strictly gains from deviating to $\bar{x}$ at those prices, so under D1 the type $\theta'$ cannot remain in the belief support. Eliminating all such types leaves only $\hat{\theta}$. If $\Sigma_T(\bar{x})$ lies outside $[\Gamma_T(0), \Gamma_T(1)]$, type monotonicity implies that the deviation gain is maximized at the corresponding endpoint, so beliefs collapse to one of the extreme types.

A key property of the commitment equilibrium is that the allocation is time-invariant. Because the individual maximizes a discounted sum of concave functions subject to the identification constraint implied by Lemma 3, optimal intertemporal smoothing emerges naturally under commitment.

LEMMA 4. *In any separating equilibrium, the allocation is time-invariant:*

$$x_t^c(\theta) = x_s^c(\theta), \quad \textit{for all } t, s \leq T \textit{ and } \theta \in [0, 1].$$

Time-invariance is central to equilibrium uniqueness. Fuchs, Gottardi, and Moreira (2025), to the contrary, find that in equilibrium quantities are indeterminate at each period, even though $\Gamma_T(\theta)$ is pinned down for each type. The difference is due to their assumption of linear payoffs for individuals and dealers.

Having established separation and time invariance under commitment, the problem reduces to the setting analyzed by Mailath (1987), yielding the following proposition.

PROPOSITION 1. *There exists a unique commitment equilibrium satisfying the D1 refinement. In this equilibrium, quantities are constant over time and evolve with type $\theta$ as the unique solution to the initial value problem*

$$x_t'(\theta) = \frac{-u_2(\theta, \theta, x_t(\theta))}{u_3(\theta, \theta, x_t(\theta))}, \text{ with } \begin{cases} x_t(1) = x^{FB}(1), \text{ if } u_2 < 0 \\ x_t(0) = x^{FB}(0), \text{ if } u_2 > 0 \end{cases} \tag{5.2}$$

*for which $x_t'(\theta)$ and $u_{1,3}$ have the same sign.*

When individuals can commit to future trades, signaling operates through quantities, not timing, along the equilibrium path. In particular, the within-period allocation is identical for any horizon $T$, a property that fails without commitment (see next section). The intuition is simple: the dynamic force in our problem is the temptation to distort today's offer to secure better prices tomorrow. Commitment removes this incentive: with no intertemporal signaling to manage, every type repeats the static allocation each period.

In the insurance example, the highest-risk type selects full insurance each period, while safer types signal by choosing lower coverage and bearing more risk. Under the specification in Section 4.1, the coverage schedule $x_t^c(\theta)$ satisfies $x_t'(\theta) = x_t(\theta)/\left(\gamma\sigma^2(1 - x_t(\theta))\right)$ with $x_t(1) = 1$.

# 6 Repeated Trade Without Commitment

When the individual cannot commit to future offers, choices are made sequentially. Past trades are publicly observable, so dealers' pricing strategies at $t$ may depend not only on the current offer $x_t$, but also on the history $h^{t-1} = \{x^{t-1}, p^{t-1}, a^{t-1}\}$, where $x^{t-1}$ is the sequence of past offers, $p^{t-1}$ the sequence of best submitted prices, and $a^{t-1}$ the sequence of acceptance decisions. Naturally, $h^{-1} = \emptyset$.

Denote dealer $i$'s pricing strategy by $p_{it}(x_t, h^{t-1})$, and write $p_t(x_t, h^{t-1})$ for the best one across dealers of period $t$. A type-$\theta$ individual's strategy is a sequence of functions $\{x_t(\theta, h^{t-1})\}_{t=0}^{T}$ mapping histories into current offers and a sequence of acceptance decisions $\{a_t(\theta, p_t, x_t, h^{t-1})\}_{t=0}^{T}$ with $a_t \in \{0, 1\}$. In a sequential equilibrium, choices must be optimal at each period, given the realized history. Let $\mathcal{U}_t(\theta, h^{t-1})$ denote the continuation value at the start of date $t$, with terminal condition $\mathcal{U}_{t+1}\left(\theta, h^{T}\right) = 0$. The dynamic problem can be written recursively as

$$\mathcal{U}_t(\theta, h^{t-1}) = \max_{x_t, a_t} w(\theta, p_t(x_t, h^{t-1}), x_t)\, a_t(\theta, p_t(x_t, h^{t-1}), x_t, h^{t-1}) + \beta \mathcal{U}_{t+1}(\theta, h^{t}). \quad (6.1)$$

Definition 2 formalizes our *adaptation* of the D1 refinement to a sequential environment. Cho and Kreps (1987) define D1 for static signaling games. In our setting the extension is natural: the beliefs being refined are those of short-lived dealers at each date, so the off-path belief restriction applies period by period. Right after the definition we discuss its sequential meaning.

In equilibrium, let $K(x_t, h^{t-1})$ denote the set of types who optimally choose $x_t$ following history $h^{t-1}$ on-path. Denote by $\Omega_t := \mathbb{R}^{t+1}$ the space of truncated contracts up to period $t$.

DEFINITION 2 (EQUILIBRIUM WITHOUT COMMITMENT). A sequential equilibrium satisfying D1 in the market without commitment is a tuple $(\{x_t^*\}, \{\mu_{it}^*\}, \{p_{it}^*\}, \{a_t^*\})$ consisting of: a sequence of contract selection strategies $x_t^*(\theta, h^{t-1})$, specifying the offer chosen by each type $\theta$ at period $t$ following history $h^{t-1}$; a system of dealer beliefs $\mu_{it}^*(x_t, h^{t-1})$, assigning probability distributions over types conditional on the observed pair $(x_t, h^{t-1})$; dealer pricing strategies $p_{it}^*(x_t, h^{t-1})$; and acceptance strategies $a_t^*(\theta, p_t, x_t, h^{t-1})$, indicating whether a type-$\theta$ individual accepts the best available price $p_t$ for offer $x_t$, such that:

(*2.1*) *Individuals optimize:* for each type $\theta$, the strategies $x_t^*, a_t^*$ satisfy (6.1) at every history $h^{t-1}$.

(*2.2*) *Dealers optimize:* given the strategies of the individual and of other dealers, each dealer $i$ submits $p_{it}^*(x_t, h^{t-1})$ to maximize expected profits given her beliefs.

(*2.3*) *On-path beliefs:* let $\Omega_t' \subseteq \Omega_t$ be the set of contracts offered on-path following history $h^{t-1}$. Then for all $x^t \in \Omega_t'$ dealers update beliefs using Bayes' rule. If a contract is selected by a measure-zero set of types, belief support is only required to lie within that set.

(*2.4*) *Off-path beliefs (D1 refinement):*

- *Individual's deviations*: for any $x^t \in \Omega_t \setminus \Omega'_t$, type $\theta$ is assigned zero posterior probability whenever, for any dealers' response $p_t$, there exists another type $\theta'$ with a strictly stronger incentive to deviate. That is, if $\theta'$ strictly prefers to deviate under any price that gives type $\theta$ only a weak incentive, then beliefs cannot assign weight to $\theta$.
- *Dealers' deviations*: if some dealer deviates at time $t$, future dealers' beliefs may place positive weight only on types in $K(x_t, h^{t-1})$

Although the logic of D1 parallels the commitment case, applying it in a sequential setting is subtler. Under commitment, a deviation is a fixed long-term contract, and its gain is that contract's value: the entire path is revealed ex ante, so dealers' beliefs - and hence prices at all dates - are embedded immediately. Without commitment, a deviation is a current offer $(x_t, h^{t-1})$ whose value equals the current payoff plus a continuation value that depends on the individual's optimal future behavior and on future dealers' D1-consistent pricing responses. Identification is therefore harder, because that continuation value varies by type and by history.

To see the logic in its simplest form, consider a horizon of two periods, $t \in \{0, 1\}$. If at $t = 1$ dealers price via a rule $p_1^*(x_1, h^0)$, then at $t = 0$ each dealer must anticipate that rule. Thus beliefs of all dealers $i \in \mathcal{I}_0$ must satisfy

$$\operatorname{supp}(\mu_{i0}^*(x_0)) \subseteq \bigcap_{p_0} \left\{ \arg\max_{\hat{\theta}} \left[ w\left(\hat{\theta}, p_0, x_0\right) + \max_{x_1} \left[ \beta w\left(\hat{\theta}, p_1^*(x_1, h^0), x_1\right) \right] - U(\hat{\theta}) \right] \right\}.$$

The intersection over $p_0$ captures that D1-consistent beliefs evaluate incentives under any feasible pricing response to $x_0$. The problem is rich because that response becomes part of $h^0$, which affects future prices $p_1^*(x_1, h^0)$ and, in turn, the optimal choice of $x_1$ by each type. Because the last period mirrors the commitment environment, only a single step of forward-looking reasoning is required.

When the horizon extends beyond two periods, the restriction acquires additional layers of recursion. At time $t$, beliefs about a deviation must anticipate how dealers at $t + 1$ would react to new offers, and those dealers, in turn, must form beliefs consistent with how the game would unfold from $t + 2$ onward. At each step, one must identify the type with the strongest deviation incentive at that stage, given that all future beliefs are themselves disciplined by D1. Formally, for any $t < T$, beliefs must satisfy

$$\operatorname{supp}\left(\mu_{it}^*\left(x_t, h^{t-1}\right)\right) \subseteq \bigcap_{p_t} \left\{ \arg\max_{\hat{\theta}} \left[ \sum_{k=0}^{t} \beta^k w\left(\hat{\theta}, p_k, x_k\right) + \beta^{t+1} \mathcal{U}_{t+1}\left(\hat{\theta}, h^t\right) - U\left(\hat{\theta}\right) \right] \right\}. \tag{6.2}$$

Here we still range over all feasible $p_t$, while prices up to $p_{t-1}$ are fixed by the realized history $h^{t-1}$. The continuation value $\mathcal{U}_{t+1}$ considers type-$\hat{\theta}$'s future optimal choices given $(x_t, h^{t-1})$ for each candidate response $p_t$.

Without commitment, dealers must treat any pair $(x_t, h^{t-1})$ with $t < T$ as potentially off-path. Even if such a pair occurs on path for some type, the individual may still deviate later. Hence beliefs at on-path histories must satisfy both Bayes' rule and D1.

The clause in Definition 2 regarding *dealers' deviations* ensures that pricing deviations do not alter the belief dynamics. Concretely, if after some $(x_t, h^{t-1})$ dealers at time $t$ believe they are facing a type $\hat{\theta}$ but they deviate in their price offer, then provided the individual rejects that offer, future dealers will not change their beliefs about the type.

As in the commitment case, competition among dealers guarantees that trades yield zero expected profits.

LEMMA 5. *In any equilibrium, any separating trade satisfies the dealers' break-even condition:*

$$\mathbb{E}\big[\pi\big(\theta, p_t, x_t\big) \mid x_t, h^{t-1}\big] = 0.$$

We introduce an infinitesimal opportunity cost of posting a nonzero quantity, which ensures that an individual proposes a nonzero quantity only if he intends to accept the best bid he anticipates. So, whenever the offer is separating, acceptance is automatic (i.e., $a_t = 1$). We can simplify the recursive optimization problem in (6.1) to

$$\mathcal{U}_t(\theta, h^{t-1}) = \max_{x_t} w\big(\theta, p(x_t, h^{t-1}), x_t\big) + \beta\, \mathcal{U}_{t+1}(\theta, h^t).$$

LEMMA 6. *In any equilibrium, every on-path nonzero trade must be fully separating.*

We provide intuition using the insurance case again. Our logic is of backward induction in the last period for which there is pooling. In period $T$ the problem collapses to a one-shot market with a fixed history, so the static argument of Lemma 2 rules out pooling in the last period. Now suppose pooling occurs, and let $t < T$ be the last period at which some types still pool. Focus on the lowest pooled type, $\theta'$, and let him shave his period-$t$ offer by an arbitrarily small amount. This deviation is either perceived as on or off-path. If this nearby offer is offered on-path by a lower type $\theta'' < \theta'$, competition delivers a discrete improvement in price at period $t$. From that point on, the deviator can copy the optimal continuation strategy of type $\theta''$ through periods $t+1, \ldots, T-1$, and then, at $T$, rebalance his quantity just enough to restore his cumulative identification to his true type (i.e., $\Gamma_T = \Sigma_T(x^T)$). This reasoning also follows the idea of Lemma 2, for the full history is fixed at the last period. Since $\theta'$ must be locally indifferent between his prescribed strategy and that of an arbitrarily close type $\theta'_+$, this plan secures him a discrete payoff gain. If instead the nearby offer was coming from a higher type, we reverse sign of the shave and apply a symmetric argument. If the small shave is off-path, D1 forces beliefs to concentrate on the endpoint for whom the deviation is instantaneously most profitable: the lowest pooled type. Because the continuation game that follows the small shave is identical for all pooled types, identification is determined solely by the instantaneous payoff change. In both the on- and off-path cases, an endpoint profits from a tiny deviation, so pooling cannot be sustained.

We impose a mild regularity condition to make the dynamic characterization tractable. We henceforth restrict attention to regular D1 equilibria in pure strategies.

DEFINITION 3 (REGULAR EQUILIBRIUM). A sequential equilibrium is *regular* if, for every history $h^{t-1}$ and every period $t \leq T$, the allocation strategy $x_t^*(\theta, h^{t-1})$ is continuous in $\theta$.

In the static benchmark ($T = 0$), Mailath (1987) shows that any separating equilibrium is regular. His argument relies on differentiability of payoffs in type and allocation. In a dynamic setting, however, the derivative of the continuation value at date $t$ depends on the differentiability of future and past equilibrium allocations. Even so, Mailath's logic extends to yield a consistency check for our regularity notion: whenever past and future allocations are continuous in a neighborhood of a given history, the incentive-compatible choice at that date must also be continuous in $\theta$.

With regularity in place, we can specify belief formation in a D1 equilibrium. Rather than repeating the full forward-looking logic at every step, we introduce a simple belief-selection device at each period, the $\Gamma$-*rule*. We then define an auxiliary equilibrium concept that embeds this rule at all dates $t \leq T$, the $\Gamma$-*criterion equilibrium*, and characterize it in Proposition 2. This construction is expositional, but serves as a transparent benchmark that we later map back to the D1-refined equilibrium in Proposition 3.

## 6.1 The $\Gamma$ Criterion

Given a sequential equilibrium $(\{x_t^*\}, \{\mu_{it}^*\}, \{p_{it}^*\}, \{a_t^*\})$, define for each $t \leq T$

$$\Gamma_t(\theta) \coloneqq \sum_{k=0}^{t} \beta^k \, g\big(x_k^*(\theta, h^{k-1}(\theta))\big),$$

where $h^{k-1}(\theta)$ is the on-path history induced by type $\theta$ up to date $k-1$. This extends $\Gamma_T(\theta)$, defined for the commitment setting, to all truncations $t \leq T$. Using $\Gamma_t$, we introduce an auxiliary equilibrium notion.

DEFINITION 4 ($\Gamma$-CRITERION EQUILIBRIUM). A sequential equilibrium satisfies the $\Gamma$-criterion if it obeys items (2.1)–(2.3) of Definition 2, and, at every date $t$, satisfies the $\Gamma$-rule:

$$\operatorname{supp}\left(\mu_{it}^*\left(x_t, h^{t-1}\right)\right) \subseteq \arg\min_{\theta \in [0,1]} \left|\Sigma_t(x_t, h^{t-1}) - \Gamma_t(\theta)\right|,$$

where $\Sigma_t(x_t, h^{t-1}) \coloneqq \sum_{k=0}^{t} \beta^k g(x_k)$.

For any pair $(x_t, h^{t-1})$, we call $\Sigma_t(x_t, h^{t-1})$ its *signature*. If the $\Gamma$-rule holds at period $t$, beliefs are formed by matching signatures to types via the schedules $\Gamma_t(\theta)$. In the insurance example, this means pricing as the type whose equilibrium discounted coverage up to $t$ equals the observed discounted coverage $\Sigma_t$.

The $\Gamma$-criterion is weaker than D1 in our setting and, by itself, may admit ill-behaved equilibria because it imposes a mathematical restriction without an underlying economic rationale.

For example, a fully pooling equilibrium may arise. Suppose $x^{FB}(\theta) = \bar{x}$ for all $\theta$, and consider strategies $x_t^*(\theta) = \bar{x}$, and prices $\bar{p}$ following $(\bar{x}, \ldots, \bar{x})$ such that $\mathbb{E}[\pi(\theta, \bar{p}, \bar{x})] = 0$. Since every $\Gamma_t(\theta)$ is constant, the $\Gamma$-rule places no restriction on off-path beliefs. In particular, dealers may assign probability one to the worst type after observing any deviation. This collection of strategies and beliefs satisfies all conditions of Definition 4.

The $\Gamma$-criterion becomes interesting once we restrict it to regular separating equilibria. In this class, each $\Gamma_t$ must be strictly monotonic and differentiable. Indeed, if at some period two distinct types trade nonzero quantities and induce the same value of $\Gamma_t$, then by Definition 4 those types are indistinguishable at that period, contradicting separation. Regularity yields continuity of $\Gamma_t$. Injectivity and continuity together imply the monotonicity and differentiability of all $\Gamma_t$. Moreover, there is a single type that solves the minimization in Definition 4, whether offers are on- or off-path, so any offer is priced as coming from a single type. Using these properties, we characterize the unique regular separating equilibrium that satisfies the $\Gamma$-criterion.

To simplify notation, given equilibrium schedules, we write $u_\ell(\theta, \theta, x_t(\theta, h^{t-1}(\theta)))$ and $g'(x_t(\theta, h^{t-1}(\theta)))$ as $u_{\ell,t}(\theta)$ and $g_t'(\theta)$ when no ambiguity arises.

PROPOSITION 2. *There exists a unique separating equilibrium satisfying the $\Gamma$-criterion in the market without commitment. In this equilibrium, quantities evolve with type $\theta$ as the unique solution to the initial value problem*

$$
\begin{aligned}
x_0'(\theta) &= \frac{-u_{2,0}(\theta)}{u_{3,0}(\theta) - u_{3,1}(\theta)\left(\frac{g_0'(\theta)}{g_1'(\theta)}\right)} \\
x_t'(\theta) &= \frac{-u_{2,t}(\theta)}{u_{3,t}(\theta) - u_{3,t+1}(\theta)\left(\frac{g_t'(\theta)}{g_{t+1}'(\theta)}\right)} + \beta^{-1}\,\frac{u_{2,t-1}(\theta)}{u_{3,t-1}(\theta)\left(\frac{g_t'(\theta)}{g_{t-1}'(\theta)}\right) - u_{3,t}(\theta)} \quad \textit{for } 0 < t < T \\
x_T'(\theta) &= \frac{-u_{2,T}(\theta)}{u_{3,T}(\theta)} + \beta^{-1}\,\frac{u_{2,T-1}(\theta)}{u_{3,T-1}(\theta)\left(\frac{g_T'(\theta)}{g_{T-1}'(\theta)}\right) - u_{3,T}(\theta)}
\end{aligned}
$$

$$
\textit{with } \begin{cases} x_t(1) = x^{FB}(1), \textit{ for all } t \leq T, \textit{ if } u_2 < 0 \\ x_t(0) = x^{FB}(0), \textit{ for all } t \leq T, \textit{ if } u_2 > 0 \end{cases}
$$

*for which $x_t'(\theta)$ and $u_{1,3}$ have the same sign for all $t$.*

To prove the proposition, formulate the type-$\theta$ problem as[1]

$$\max_x \sum_{t=0}^{T} \beta^t u(\theta, \hat{\theta}(x_t, h^{t-1}), x_t) \text{ subject to the } \Gamma\text{-rule at each } t.$$

Take first-order conditions and use backward substitution to obtain the system of differential equations. The second-order condition requires each $\Gamma_t(\theta)$ to be strictly increasing, and a single-crossing argument then implies that equilibrium allocations are increasing in type when $u_{1,3} > 0$,

[1] The formulation mirrors Kyle (1985), where an informed trader chooses quantities that both maximize current payoff and shape future beliefs. The key difference is that noise traders prevent perfect separation in Kyle's setting, yielding more complex equilibrium schedules than in our environment.

and decreasing otherwise. The boundary condition is pinned down by the worst type, who cannot be credibly punished for deviating and thus trades at his first-best quantities. Efficiency at the worst extreme implies that the derivatives $x_t'(\theta)$ blow up close to that point, so usual existence and uniqueness results do not apply. The Appendix proves existence and uniqueness of an equilibrium despite the singular initial condition. We analyze the system's dynamics and its interpretation in the next subsection.

## 6.2 The D1 Equilibrium

Our main result establishes an equivalence between the $\Gamma$-criterion and D1. Hence, the tractable system in Proposition 2 characterizes the seemingly complex D1 equilibrium.

PROPOSITION 3. *A regular separating equilibrium in the market without commitment satisfies D1 if and only if it satisfies the* $\Gamma$*-criterion.*

*Proof.* We first show that D1 implies the $\Gamma$-criterion. Fix any $(x_t, h^{t-1})$. Using the functional form of $w$, rewrite (6.2) as

$$\operatorname{supp}\left(\mu_{it}\left(x_t, h^{t-1}\right)\right) \subseteq \bigcap_{p_t} \left\{ \arg\max_{\hat{\theta}} \left\{ \hat{\theta}\, \Sigma_t\left(x_t, h^{t-1}\right) + \mathcal{U}_{t+1}\left(\hat{\theta}, h^t\right) - U\left(\hat{\theta}\right) \right\} \right\}. \quad (6.3)$$

Fix a price response $p_t$ and let $(\tilde{x}_{t+1}(\hat{\theta}), \ldots, \tilde{x}_T(\hat{\theta}))$ be the optimal continuation for type $\hat{\theta}$ following $(x_t, h^{t-1})$. By the envelope condition, any interior maximizer $\hat{\theta}$ in (6.3) must satisfy

$$\Gamma_T\left(\hat{\theta}\right) = \Sigma_t\left(x_t, h^{t-1}\right) + \sum_{k=t+1}^{T} \beta^k g\left(\tilde{x}_k\left(\hat{\theta}\right)\right). \quad (6.4)$$

To prove that D1 requires the $\Gamma$-rule to hold at every $t \leq T$, we argue that in any regular separating equilibrium satisfying D1, each $\Gamma_t$ is strictly monotonic and equation (6.4) admits a unique solution that coincides with the $\Gamma-$rule candidate,

$$\hat{\theta}\left(x_t, h^{t-1}\right) = \Gamma_t^{-1}\left(\Sigma_t\left(x_t, h^{t-1}\right)\right),$$

whenever $\Sigma_t(x_t, h^{t-1})$ is in the image of $\Gamma_t$, denoted by $\Gamma_t[0,1]$.

Proceed by backward induction. The base case $t = T$ reduces to the commitment case (Lemma 3). For the inductive step, assume the $\Gamma$–rule holds for all $k > t$. Suppose, toward a contradiction, that there exists an off-path history $(\bar{x}_t, \bar{h}^{t-1})$ with $\Sigma_t(\bar{x}_t, \bar{h}^{t-1}) \in \Gamma_t[0,1]$ and beliefs assigning positive probability to some $\bar{\theta}$ with $\Gamma_t(\bar{\theta}) \neq \Sigma_t(\bar{x}_t, \bar{h}^{t-1})$. Under D1, there must be a continuation $(\bar{x}_{t+1}, \ldots, \bar{x}_T)$ that makes $\bar{\theta}$ the most willing to deviate, with no alternative deviation relatively more attractive for any other type. Dynamic consistency and the inductive hypothesis then require future beliefs and prices to follow the $\Gamma$-rule, so $\Sigma_k(\bar{x}_k, \bar{h}^{k-1}) = \Gamma_k(\bar{\theta})$ for all $k > t$. By sequential optimality for type $\bar{\theta}$ on path, this continuation cannot be optimal unless $\Sigma_t = \Gamma_t(\bar{\theta})$ (we formalize this claim in Appendix A.7), a contradiction.

The argument does not depend on the particular price response $p_t$, so the conclusion holds for all feasible $p_t$, and the intersection in (6.3) contains a unique $\hat{\theta}$, namely $\Gamma_t^{-1}(\Sigma_t)$. Considering the case $\Sigma_t \in \Gamma_t[0,1]$ suffices for on-path characterization. Once the equilibrium is constructed, when $\Sigma_t \notin \Gamma_t[0,1]$, D1 assigns the extreme types 0 (when $\Sigma_t < \Gamma_t(0)$) or 1 (when $\Sigma_t > \Gamma_t(1)$), which is consistent with the $\Gamma$-rule (Definition 4).

For the converse, Proposition 2 constructs an equilibrium under the $\Gamma$-rule in which off-path offers are priced by the most-willing type, and the on-path allocation arises from each type's optimization over correctly priced deviations. Hence the $\Gamma$-criterion implies D1. □

We now discuss some results on equilibrium dynamics. Figure 1 illustrates them by plotting equilibrium coverages in the insurance example for $T = 1$ and $T = 25$.

Figure 1 – Equilibrium Coverages in the Insurance Example

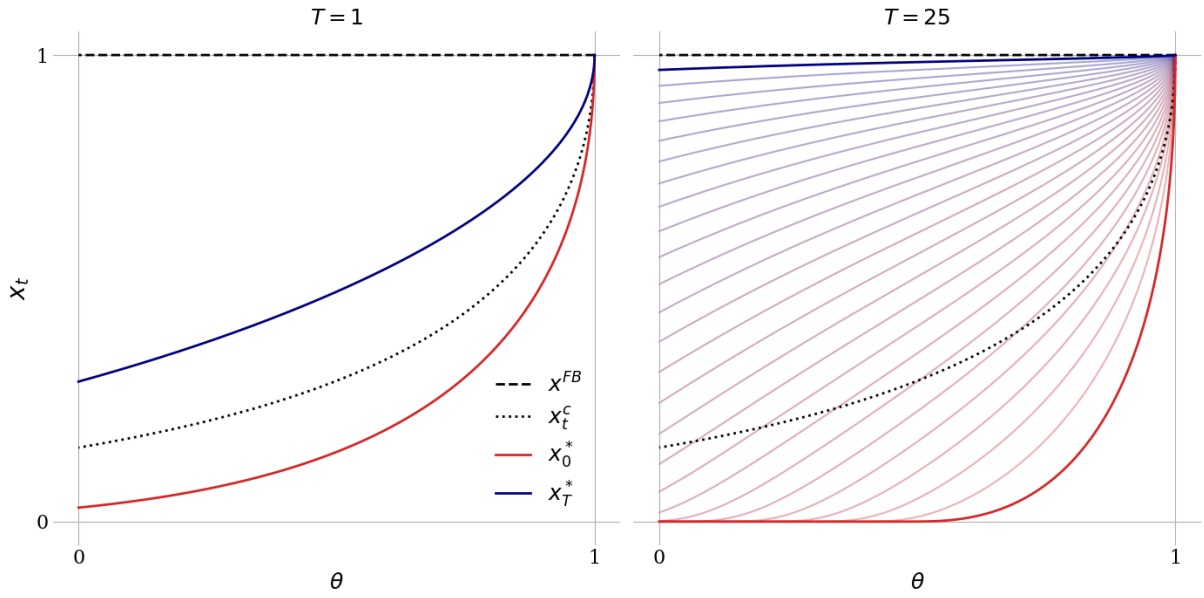


*Notes:* Parameters are set to $\gamma = \sigma^2 = 1$ and $\beta = 0.9$. The dashed black $x^{FB}(\theta) = 1$ is the stage first-best, and the dotted black line $x_t^c(\theta)$ is the commitment schedule (time-invariant in $t$). Without commitment, red is the initial schedule $x_0^*(\theta)$ and blue is the terminal schedule $x_T^*(\theta)$. The left panel reports $T = 1$. The right panel reports $T = 25$, where intermediate allocations shade from red to blue as $t$ increases.

COROLLARY 4. *The equilibrium allocation for each type evolves monotonically toward $x^{FB}(\theta)$. For all $t < T$ and all $\theta \in (0,1)$ it holds that $x_t(\theta) < x_{t+1}(\theta)$ when type and belief monotonicity have opposite signs; and $x_t(\theta) > x_{t+1}(\theta)$ when they have the same sign.*

In the insurance application, coverage rises monotonically over time toward full-insurance for all $\theta < 1$. This follows directly from the equilibrium system and the fact that each $\Gamma_t(\theta)$ is increasing. Moreover, since $\beta^T$ vanishes as $T$ grows, terminal-period allocations get arbitrarily close to the efficient ones as we allow for more periods in the game. Otherwise $\Gamma_T$ would approach zero, contradicting separation. This yields the next result.

COROLLARY 5. *When $T \to \infty$, the trades $x_t(\theta)$ approach their first-best quantities $x^{FB}(\theta)$ as $t \to T$.*

Taken together, the corollaries imply that information frictions unwind over time. In early periods, dynamic incentives amplify the static distortion: anticipating that worse types may mimic them to secure better prices, all types (except for the worst one) scale back quantities relative to the static benchmark. As history accumulates and those early signaling costs are sunk and observed, the marginal value of further signaling falls, distortions shrink, and allocations move monotonically toward the first-best.

There are many ways to rearrange the system in Proposition 2. In particular, we can express it in terms of $\Gamma_t'(\theta)$. For simplicity, and as in the insurance application, consider $g(x_t) = x_t$. The system then becomes:

$$\Gamma_t'(\theta) = \beta^t \frac{-u_{2,t}}{u_{3,t} - u_{3,t+1}} \text{ for } t < T, \text{ and } \Gamma_T'(\theta) = \beta^T \frac{-u_{2,T}}{u_{3,T}}.$$

This is directly comparable to the commitment equilibrium in Proposition 1. With dynamic incentives, current marginal changes in the on-path signatures $\Gamma_t(\theta)$ equal the discounted ratio of current informational costs $(-u_2)$ to the difference in marginal gains from quantities between the current and next periods $(u_{3,t} - u_{3,t+1})$. The expression for $\Gamma_t'(\theta)$ shows how expected future retention raises the cost of current separation by increasing $u_{3,t+1}$. As history accumulates and quantities approach the efficient benchmark, separation becomes progressively less costly. Hence repeated trading endogenously weakens the signaling role of retention over time.

An alternative interpretation of the system in Proposition 2 is to view the derivatives $x_t'(\theta)$ directly as the sum of two terms. The first term relates marginal gains from trade at dates $t$ and $t+1$ (as reflected in the denominators). Anticipated increases in future quantities raise the need for distortions today. The second term has the opposite sign and captures relative marginal gains at $t-1$ and $t$. This term shows how accumulated past history, summarized solely by yesterday's trade, allows for lower distortions in current quantities.

Heuristically, quantity distortions are smaller when utility is more concave in quantities and when discounting is stronger. Greater curvature or a lower discount factor raises the marginal cost of current deviations, strengthens the signal, and reduces the efficiency wedge. We provide clean comparative statics for the insurance case in Section 7.1.

Timing matters: both the interval between trading dates and the game horizon directly affect the equilibrium retention. Changes in the interval between dates are equivalent to changing the individual's discount factor. So, for any horizon, distortions are greater each period when trading dates are closer to each other. The same happens with an increase in the game horizon, $T$. The more trading opportunities available until the end of the game, the greater the individual's incentive to distort current quantities, as he will enjoy the benefits of better price schedules for a higher number of periods ahead.

## 6.3 Observability

Transparency of past trades is pivotal to our analysis. If past trades are not publicly observed, two forces collapse the problem into a static one each period. First, with no history to condition on, beliefs reset each date, and each period is effectively a fresh start. Second, with nothing to infer from unobserved behavior, today's choice cannot influence future pricing, so each period also behaves like the last. Hence, the individual and dealers effectively play the same one-shot game in every period. Because the commitment allocation is time-invariant and independent of the horizon $T$, the commitment schedule is implemented without commitment when past trades are opaque.

This parallel between opacity and commitment has interesting welfare implications. If a policymaker could choose ex ante between opaque and transparent regimes, opacity would be preferred:

LEMMA 7. *All types weakly prefer the equilibrium sequence of contracts under commitment to that under no commitment.*

However, when there is no commitment and trading histories are observable, allocations converge to efficiency over time. For sufficiently long horizons there exists a date $\tau$ after which discounted continuation welfare under transparency strictly exceeds that under opacity. Figure 2 plots continuation welfare, as a share of the efficient benchmark, under both regimes with uniform weights.

Figure 2 – Continuation Social Welfare (% of First-Best)

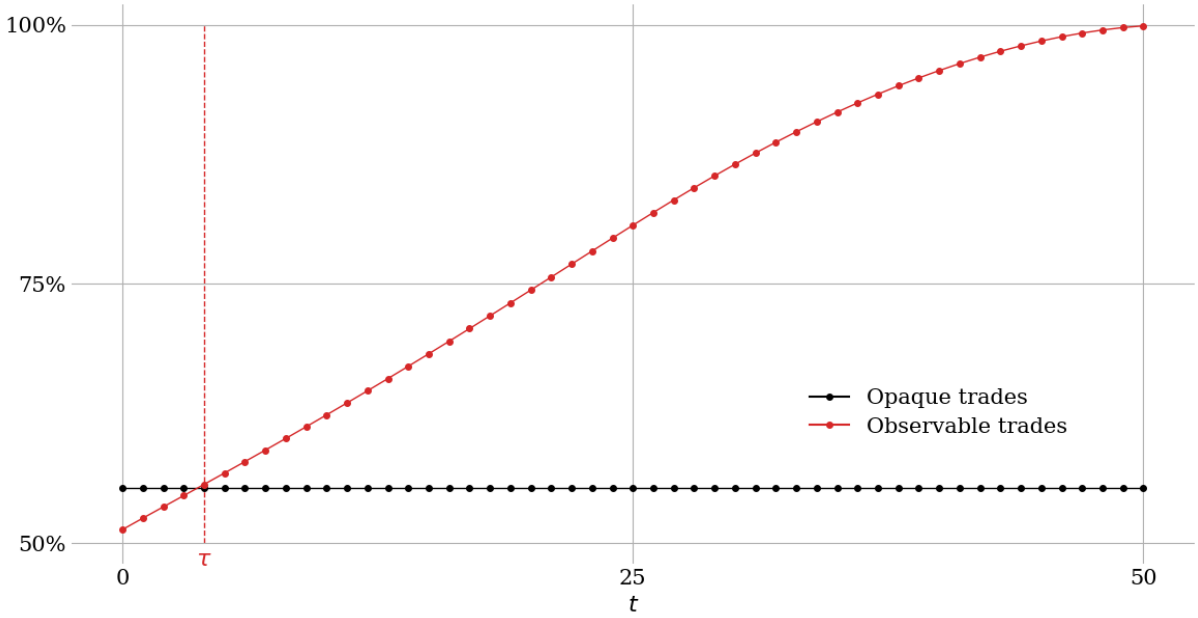


*Notes:* Dynamic insurance specification with $\gamma = \sigma^2 = 1$ and $\beta = 0.99$. Points report, at each date $t$, discounted continuation welfare $\sum_{k=t}^{T} \beta^{k-t} \int_0^1 u(\theta, \theta, x_k^*(\theta)) d\theta$ relative to the first-best benchmark.

Observability creates an intertemporal trade-off: early periods bear high signaling costs, while later periods reap the payoff from those early distortions. In welfare terms, this yields a front-loaded loss relative to the opaque benchmark and a back-loaded gain when the horizon is sufficiently long. Policy implications are thus path-dependent. If public observability has been in place long enough that initial distortion costs are sunk, maintaining transparency is socially desirable. If the market is opaque, it is better to keep it opaque than to bear the initial costs of introducing observability.

A welfare loss from transparency is consistent with prior results (e.g., Hörner and Vieille (2009) and Fuchs, Öry, and Skrzypacz (2016)). Relatedly, Kaya and Roy (2024) show in a two-type model of repeated trading that transparency can raise or lower welfare depending on dealers' priors. The key difference is that their good is indivisible, so trades need not separate types and dynamics hinge on initial beliefs, whereas our divisible good environment delivers separation along the path.

# 7 Applications: Results

We turn back to the applications of Section 4 and discuss how our results apply to each.

## 7.1 Insurance

In repeated insurance markets, low-risk individuals initially distort coverage downward. By retaining more risk early on, they build a credible track record that signals their quality, which then allows them to raise coverage over time while still obtaining fair prices (see Figure 1).

Under the CARA-Gaussian specification, we conduct comparative statics with respect to the discount factor $\beta$, risk aversion $\gamma$, and shock variance $\sigma^2$.

LEMMA 8. *For all $t \leq T$ and $\theta \in [0,1]$, lower discount factors, higher risk aversion, and higher shock variance all, ceteris paribus, increase equilibrium coverage:*

- *if $\beta_0 < \beta_1$, then $x_t^*(\theta \mid \beta_0) \geq x_t^*(\theta \mid \beta_1)$;*
- *if $\gamma_1 > \gamma_0$, then $x_t^*(\theta \mid \gamma_1) \geq x_t^*(\theta \mid \gamma_0)$;*
- *if $\sigma_1^2 > \sigma_0^2$, then $x_t^*(\theta \mid \sigma_1^2) \geq x_t^*(\theta \mid \sigma_0^2)$.*

To prove the lemma, we incorporate the parameters as initial conditions in the system and apply the Kamke monotonicity condition for dynamical systems (Smith, 1995). The intuition for the comparative statics is as follows. For $\beta$ at $t = 0$, a lower discount factor reduces the value of future price improvements, weakening the incentive to distort coverage today to deter mimicry tomorrow, so $x_0^*$ rises. Higher past coverage relaxes current signaling needs, pushing $x_t^*$ up for all $t > 0$ when $\beta$ falls. For $\gamma$, higher risk-aversion raises the marginal cost of retaining risk, making each unit of under-coverage a stronger signal. Less distortion is needed, so coverage increases. The same logic applies to $\sigma^2$. These comparisons are in terms of equilibrium quantities. Welfare effects need not be monotone in $\beta, \gamma, \sigma^2$ because the value functions themselves depend on these parameters.

## 7.2 Financing

We return to the case of selling a divisible asset to risk-averse buyers. The period-$t$ signature is

$$\Sigma_t(x_t, h^{t-1}) = -\sum_{k=0}^{t} \beta^k x_k,$$

the negative of total discounted sales up to $t$. Under the $\Gamma$-rule, buyers at date $t$ form beliefs from this cumulative quantity (including the current offer). Type monotonicity is negative, so higher-quality types signal by selling less each period. Belief monotonicity is positive, so the boundary condition is pinned down by type $0$. Without commitment, all $\theta > 0$ undersell

initially relative to first-best and then increase sales over time, converging to efficiency. Figure 3 illustrates these dynamics. They speak immediately to the retention logic of Leland and Pyle (1977), and exemplify our result that repeated trading endogenously weakens the signaling role of retention over time.

A related strand on credit rationing and relationship lending reaches similar conclusions.[1] Sharpe (1990) presents a model where capital allocation is shifted toward riskier types and, in equilibrium, reputation improves efficiency over time. Empirically, Petersen and Rajan (1994) find that relationships increase financing availability over time and operate primarily through quantities rather than prices, which is perfectly consistent with our model. Our work fits in this literature once we reinterpret the establishment of a public history as a relationship between the firm and the market.

Figure 3 – Equilibrium Quantities in the Financing Application

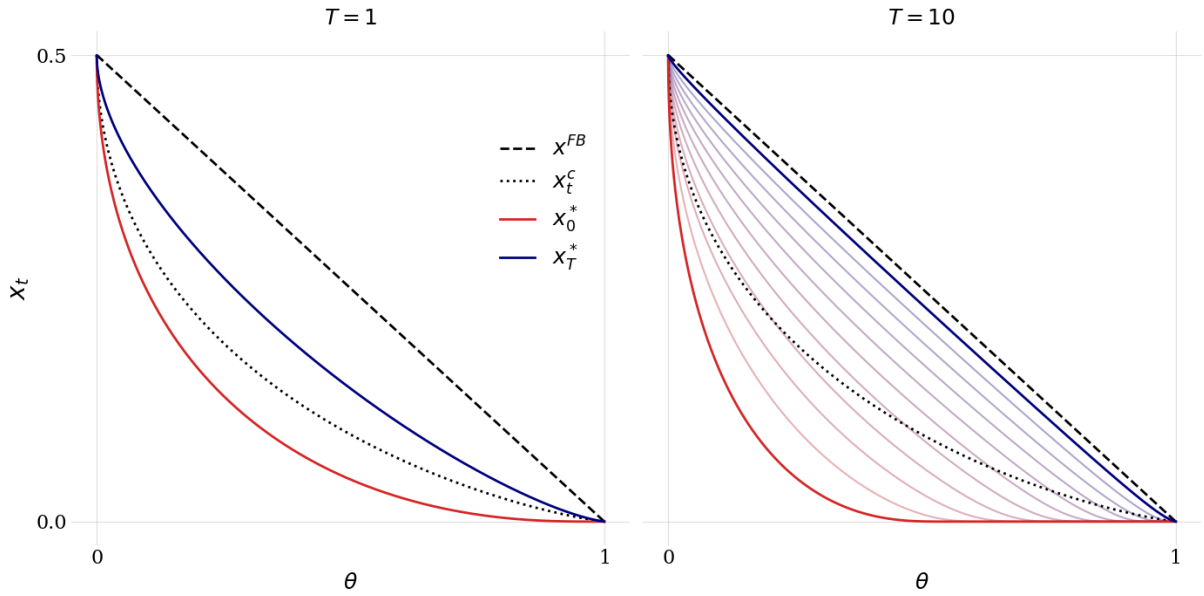


*Notes:* $v(\theta) = (1+\theta)/2$, $\gamma = \sigma^2 = 1$, $\beta = 0.9$. The dashed black line is the stage first-best $x^{FB}(\theta)$. The dotted black line is the (time-invariant) commitment schedule $x_t^c(\theta)$. Without commitment, red is $x_0^*(\theta)$ and blue is $x_T^*(\theta)$. The left panel reports $T = 1$ and the right panel reports $T = 10$, where intermediate allocations shade from red to blue as $t$ increases.

## 7.3 Job Market

In the job-market case, where workers can contract any number of hours each period, equilibrium depends on how work costs vary with type. Consider the case in which higher-quality workers dislike working more, consistent with higher opportunity costs. If $c(\theta, x_t) = \theta x_t + x_t^2$, the period-$t$ signature is the negative of total discounted hours up to $t$. High types then signal by working less, echoing the countersignaling logic in Feltovich, Harbaugh, and To (2002) and Araujo, Gottlieb, and Moreira (2007). As history accumulates and high types have "shown

[1] See Fried and Howitt (1980) on theoretical links among credit, insurance and labor contracts.

confidence" by working suboptimally few hours early on, they can raise hours offered and move toward their efficient working hours.

Our analysis also offers a complementary mechanism for inefficient hiring in entry-level labor markets. While Pallais (2014) attributes initial mismatch to limited information, so that efficiency improves as firms learn from on-the-job performance, we show that inefficient early hiring can also arise endogenously as a signaling device, depending on how work-related costs correlate with worker ability.

Alternatively, if higher types face lower work costs ($c_{1,2}(\theta, x_t) < 0$), Assumption 2 would not hold. However, under stronger conditions on $c$, we could still obtain equilibrium uniqueness. The period-$t$ signature then equals the (positive) total discounted hours up to $t$. High types initially work more than the efficient quantity and then gradually scale back as their careers progress. This pattern connects to the empirical evidence in Anger (2008), who finds that overtime can serve as a signal to employers.

# 8 Extension: Capacity Constraint

In many settings, past trades shape not only dealers' beliefs but also the individual's future feasibility set. A natural example is our financing application when the seller owns a single divisible unit to sell over time.

The characterizations in Sections 5 and 6 relied on interior choices at every date. With a capacity constraint, off-path belief formation becomes more involved: buyers must evaluate optimal continuations for each type at each residual supply level, and these continuations often hit corners. A general $T$-period analysis is therefore cumbersome. We focus on a two-period model to isolate how capacity changes the logic.

We adapt the finance setting in Section 4.2 to a single-unit supply. The adapted model is very close to that of Fuchs, Gottardi, and Moreira (2025), except that buyers' payoffs are concave in quantities rather than linear. We then characterize equilibria with and without commitment. Some properties parallel theirs, others mirror our previous analysis.

## 8.1 Adapted Model

There are two trading dates, $t \in \{0, 1\}$. A type-$\theta$ seller owns one unit of a durable, divisible asset. His payoff from $(p_t, x_t)$ is $w(\theta, p_t, x_t) = p_t - \theta\, x_t$. Dealer $i$, if she pays $p_{it}$ for $x_t$ units of quality $\theta$, earns $\pi(\theta, p_{it}, x_t) = v(\theta)\, \phi(x_t) - p_{it}$, where $v$ is a valuation function similar to that of Section 4.2, $\phi$ is a strictly increasing and strictly concave function with $\phi(0) = 0$, and $v(\theta)\phi'(x_t) - \theta$ is always positive.

Assumption 1 does not hold here because the analogue of the finite-solution condition 1.C fails: the stage problem has no interior maximizer. When constrained, the seller wishes to sell his full unit by the end of the game regardless of his type, and would sell more if he owned more than one unit. There is no meaningful *stage* first-best, the first-best is an intertemporal split of the unit, which features decreasing offers over time for all types because of the discount factor. Equilibrium notions with and without commitment are as in Definitions 1 and 2, with the added feasibility constraint $x_0 + x_1 \leq 1$. In this setting, continuity of equilibrium allocation arises endogenously, so we do not need to focus on regular equilibria.

## 8.2 Commitment

Under commitment, Lemmas 1-3 continue to hold: competition yields zero profits, the equilibrium is separating, and beliefs are pinned down by the signature $\Sigma_1(x) = -(x_0 + \beta x_1)$. Because there is no distortion at the bottom, type 0 front-loads, so $x_0^c(0) > x_1^c(0)$, and time invariance fails. Nonetheless, the logic behind Lemma 4 still applies: sellers maximize utility subject to their identification constraint. By concavity and continuity of $u$, types near $\theta = 0$ also sell the full unit but rebalance across periods. As $\theta$ rises there exists $\bar{\theta}^c$ with $x_0^c(\bar{\theta}^c) = x_1^c(\bar{\theta}^c) = \frac{1}{2}$.

If $\bar{\theta}^c < 1$, then for $\theta \in [\bar{\theta}^c, 1]$ the allocation is time-invariant, the capacity constraint is slack, and the equilibrium in this interval mirrors the commitment analysis in Section 5.

## 8.3 No Commitment

Without commitment, belief formation under D1 follows the logic of Section 6, subject to the feasibility constraint $x_0 + x_1 \leq 1$. The equilibrium has two phases: a binding-capacity phase and a slack-capacity phase. The intuition is clearest via a thought experiment. Fix a capacity constrained no-commitment equilibrium. Now relax the capacity constraint for a single type $\tilde{\theta}$, keeping all other types constrained, and ask whether $\tilde{\theta}$ would profitably deviate.

Type $\theta = 0$ would surely deviate if unconstrained. Indeed, he sells according to first-best on the equilibrium path, so his capacity constraint binds and no credible punishment could deter an increase. By continuity, the same is true for types $\tilde{\theta}$ close to $0$. Thus, there exists a threshold $\bar{\theta}^* > 0$, the lowest type who would not deviate if he alone were unconstrained. At $\bar{\theta}^*$, the marginal gain from raising $x_1^*(\bar{\theta}^*)$ is exactly offset by the marginal loss from the induced change in valuation. Consequently, for all types in the first phase, $[0, \bar{\theta}^*]$, the capacity constraint binds at the terminal date, whereas for types in the second phase, $(\bar{\theta}^*, 1]$, there is slack. In Appendix A.10 we show that $\bar{\theta}^c < \bar{\theta}^*$.

The equilibrium construction proceeds in two steps. First, solve the binding-capacity phase. For low types, dealers at $t = 0$ anticipate that the entire residual supply will be sold at $t = 1$ and form beliefs accordingly. This delivers the first phase schedule and pins down $\bar{\theta}^*$. Second, on the slack-capacity phase, apply the characterization in Propositions 2 and 3, with initial conditions at $\bar{\theta}^*$ matching the first phase values $\big(x_0^*(\bar{\theta}^*), x_1^*(\bar{\theta}^*)\big)$, and dropping the requirement that $x_t'(\theta)$ be monotonic. Figure 4 illustrates the equilibrium quantities for both commitment and no commitment with $\phi(x_t) = x_t^\alpha/\alpha$ and $v(\theta) = (1 + \theta)/2$.

With two trading dates, dealers at $t = 0$ correctly anticipate future trades, and the seller cannot benefit from deviating later. Consequently, on $[0, \bar{\theta}^c]$, the no-commitment allocation coincides with the commitment allocation. Thus, wherever the commitment outcome sells the full unit, it is implementable without commitment, echoing the main result in Fuchs, Gottardi, and Moreira (2025). Moreover, on $[0, \bar{\theta}^*]$ our no-commitment equilibrium shares their feature that total sales exhaust the residual supply whenever trade begins before the final date. The second phase, $(\bar{\theta}^*, 1]$, departs from their characterization due to curvature and fits naturally into the framework developed in Section 6.

In the two-period case the $\Gamma-$rule used to construct the second phase also applies trivially in the first phase, but this equivalence fails in general for $T > 1$. For longer horizons, the equilibrium retains the two-phase structure, and explicit characterizations can be obtained for small $T$. However, belief formation in the binding-capacity phase becomes increasingly intricate due to corner solutions, so the first phase system does not generalize cleanly.

Figure 4 – Equilibrium Quantities by Type with Capacity Constraint

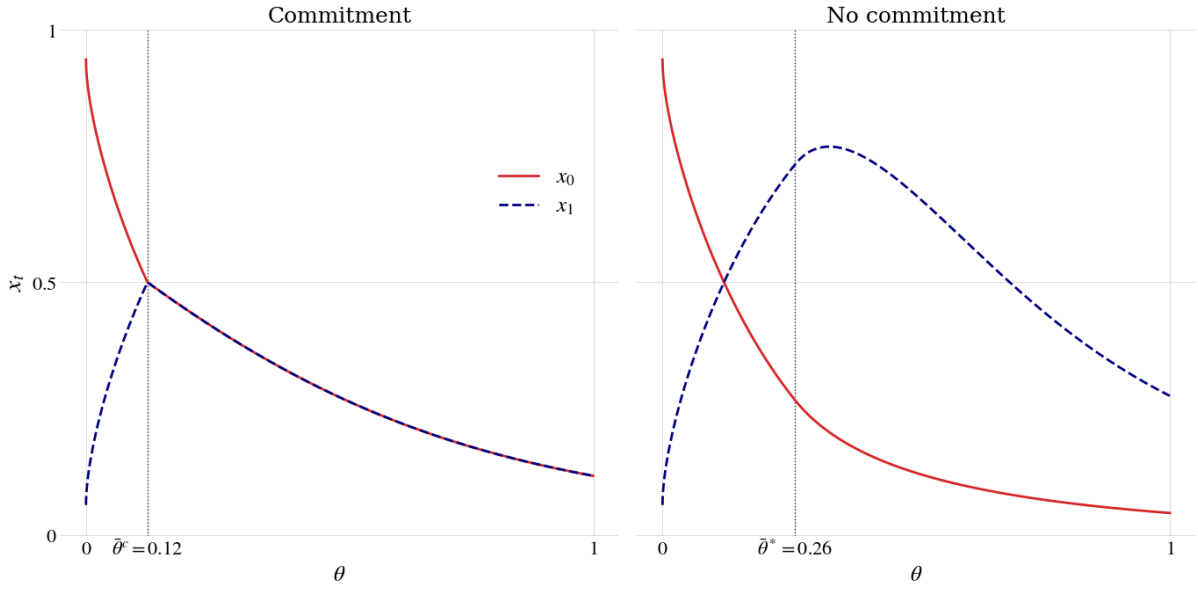


*Notes:* $v(\theta) = (1+\theta)/2$, $\phi(x_t) = x_t^\alpha/\alpha$ with $\alpha = 0.75$, and $\beta = 0.5$. Red: $t = 0$. Blue dashed: $t = 1$. Left panel shows the commitment case, with transition at $\bar{\theta} = 0.12$. Right panel shows the no commitment case, with transition $\bar{\theta} = 0.26$.

# 9 Conclusion

We analyzed adverse selection in a repeated-trading environment with a divisible good or service: a long-lived, privately informed individual trades over time with short-lived competitive dealers. We have introduced a belief formation rule (the $\Gamma$-rule) that maps observed trade histories into types via a single-dimensional statistic, and have shown that, under regularity and separation, it is equivalent to D1. The equivalence collapses a high-dimensional dynamic signaling problem into a tractable system of first-order differential equations. In the commitment benchmark, we obtain separation and intertemporal smoothing. Without commitment, separation is also achieved, and we find a unique and history-dependent equilibrium in which quantities move monotonically toward the stage first-best and converge to efficiency as the horizon lengthens. Applications to dynamic insurance, financing with risk-averse buyers, and the job market illustrate how signals accumulate and distortions unwind over time. A capacity-constraint extension highlights how corners complicate off-path beliefs and produce a two-phase equilibrium. We also discussed transparency of the trading history: opacity replicates commitment ex ante, whereas observability imposes front-loaded signaling costs but delivers back-loaded efficiency gains, leading to path-dependent welfare conclusions.

Our modeling choice of linear-in-type payoffs and our analytical focus on regular D1 equilibria may narrow the range of immediate applications, but they yield sharp and transparent results. Relaxing these features would broaden the scope, while posing substantial technical challenges.

We highlight three promising directions for future research. First, allow only partial or noisy observability of past trades. Second, move beyond a single trait to multi-dimensional or partly stochastic private information. These two directions raise common questions: when beliefs collapse to an effectively one-dimensional statistic, when bunching arises, whether there is convergence to the efficient allocation over time, and, if so, how the speed of convergence scales with noise and cross-dimensional correlation. A third direction is normative: design disclosure regimes, record-keeping rules, or simple commitment devices and quantify the trade-off between transitional signaling costs and dynamic efficiency gains. Recent work has advanced these questions, but typically without the divisibility that is central in our setting. Our work suggests that incorporating divisibility should be especially informative.

# References


Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*(3), 488–500.

Anger, S. (2008). Overtime work as a signaling device. *Scottish Journal of Political Economy 55*(2), 167–189.

Araujo, A., D. Gottlieb, and H. Moreira (2007). A model of mixed signals with applications to countersignalling. *The RAND Journal of Economics 38*(4), 1020–1043.

Asriyan, V. and V. Vanasco (2024). Security design in non-exclusive markets with asymmetric information. *Review of Economic Studies 91*(2), 690–719.

Attar, A., T. Mariotti, and F. Salanié (2011). Nonexclusive competition in the market for lemons. *Econometrica 79*(6), 1869–1918.

Azevedo, E. M. and D. Gottlieb (2017). Perfect competition in markets with adverse selection. *Econometrica 85*(1), 67–105.

Biais, B., D. Martimort, and J.-C. Rochet (2000). Competing mechanisms in a common value environment. *Econometrica 68*(4), 799–837.

Bisin, A. and P. Gottardi (2006). Efficient competitive equilibria with adverse selection. *Journal of Political Economy 114*(3), 485–516.

Blair, P. Q., T. G. Castagnino, E. L. Groshen, P. Debroy, B. Auguste, S. Ahmed, F. G. Diaz, and C. Bonavida (2020). Searching for stars: Work experience as a job market signal for workers without bachelor's degrees. Technical report, National Bureau of Economic Research.

Camargo, B. and B. Lester (2014). Trading dynamics in decentralized markets with adverse selection. *Journal of Economic Theory 153*, 534–568.

Cho, I.-K. and D. M. Kreps (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics 102*(2), 179–221.

Courty, P. and L. Hao (2000). Sequential screening. *The Review of Economic Studies 67*(4), 697–717.

Daily, G., I. Hendel, and A. Lizzeri (2008). Does the secondary life insurance market threaten dynamic insurance? *American Economic Review 98*(2), 151–156.

Daley, B. and B. Green (2012). Waiting for news in the market for lemons. *Econometrica 80*(4), 1433–1504.

DeMarzo, P. and D. Duffie (1999). A liquidity-based model of security design. *Econometrica 67*(1), 65–99.

Deneckere, R. and M.-Y. Liang (2006). Bargaining with interdependent values. *Econometrica 74*(5), 1309–1364.

Dilmé, F. (2025). Repeated trade with imperfect information about previous transactions. *Theoretical Economics 20*(1), 209–254.

Dionne, G. (1981). Adverse selection and repeated insurance contracts. In *Foundations of Insurance Economics: Readings in Economics and Finance*, pp. 407–423. Springer.

Dionne, G. and N. A. Doherty (1994). Adverse selection, commitment, and renegotiation: Extension to and evidence from insurance markets. *Journal of Political Economy 102*(2), 209–235.

Du, S. and H. Zhu (2017). What is the optimal trading frequency in financial markets? *The Review of Economic Studies 84*(4), 1606–1651.

Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013). Selection on moral hazard in health insurance. *American Economic Review 103*(1), 178–219.

Eső, P. and B. Szentes (2007). Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies 74*(3), 705–731.

Evans, R. (1989). Sequential bargaining with correlated values. *The Review of Economic Studies 56*(4), 499–510.

Farinha Luz, V., P. Gottardi, and H. Moreira (2023). Risk classification in insurance markets with risk and preference heterogeneity. *Review of Economic Studies 90*(6), 3022–3082.

Feltovich, N., R. Harbaugh, and T. To (2002). Too cool for school? signalling and countersignalling. *RAND Journal of Economics 33*(4), 630–649.

Fried, J. and P. Howitt (1980). Credit rationing and implicit contract theory. *Journal of Money, Credit and Banking 12*(3), 471–487.

Fuchs, W., P. Gottardi, and H. Moreira (2025). Time trumps quantity in the market for lemons. Working Paper.

Fuchs, W., A. Öry, and A. Skrzypacz (2016). Transparency and distressed sales under asymmetric information. *Theoretical Economics 11*(3), 1103–1144.

Fuchs, W. and A. Skrzypacz (2015). Government interventions in a dynamic market with adverse selection. *Journal of Economic Theory 158*, 371–406.

Fuchs, W. and A. Skrzypacz (2019). Costs and benefits of dynamic trading in a lemons market. *Review of Economic Dynamics 33*, 105–127.

Gerardi, D., L. Maestri, and I. Monzón (2022). Bargaining over a divisible good in the market for lemons. *American Economic Review 112*(5), 1591–1620.

Guerrieri, V., R. Shimer, and R. Wright (2010). Adverse selection in competitive search equilibrium. *Econometrica 78*(6), 1823–1862.

Hart, O. D. and J. Tirole (1988). Contract renegotiation and coasian dynamics. *The Review of Economic Studies 55*(4), 509–540.

Hendel, I. and A. Lizzeri (2003). The role of commitment in dynamic contracts: Evidence from life insurance. *The Quarterly Journal of Economics 118*(1), 299–328.

Hörner, J. and N. Vieille (2009). Public vs. private offers in the market for lemons. *Econometrica 77*(1), 29–69.

Hosios, A. J. and M. Peters (1989). Repeated insurance contracts with adverse selection and limited commitment. *The Quarterly Journal of Economics 104*(2), 229–253.

Janssen, M. and S. Roy (2002). Dynamic trading in a durable good market with asymmetric information. *International Economic Review 43*(1), 257–282.

Kaya, A. and K. Kim (2018). Trading dynamics with private buyer signals in the market for lemons. *The Review of Economic Studies 85*(4), 2318–2352.

Kaya, A. and S. Roy (2024). Repeated trading: Transparency and market structure. *American Economic Review 114*(8), 2388–2435.

Kremer, I. and A. Skrzypacz (2007). Dynamic signaling and market breakdown. *Journal of Economic Theory 133*(1), 58–82.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica 53*(6), 1315–1335.

Leland, H. E. and D. H. Pyle (1977). Informational asymmetries, financial structure, and financial intermediation. *The Journal of Finance 32*(2), 371–387.

Levy, Y. J. and A. Veiga (2025). Optimal contract regulation in selection markets. *American Economic Journal: Microeconomics 17*(2), 94–126.

Li, Q. (2022). Security design without verifiable retention. *Journal of Economic Theory 200*, 105381.

Mailath, G. J. (1987). Incentive compatibility in signaling games with a continuum of types. *Econometrica 55*(6), 1349–1365.

Myers, S. C. and N. S. Majluf (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics 13*(2), 187–221.

Noldeke, G. and E. Van Damme (1990). Signalling in a dynamic labour market. *The Review of Economic Studies 57*(1), 1–23.

Pallais, A. (2014). Inefficient hiring in entry-level labor markets. *American Economic Review 104*(11), 3565–3599.

Pavan, A., I. Segal, and J. Toikka (2014). Dynamic mechanism design: A myersonian approach. *Econometrica 82*(2), 601–653.

Petersen, M. A. and R. G. Rajan (1994). The benefits of lending relationships: Evidence from small business data. *The Journal of Finance 49*(1), 3–37.

Ramey, G. (1996). D1 signaling equilibria with multiple signals and a continuum of types. *Journal of Economic Theory 69*(2), 508–531.

Rey, P. and B. Salanie (1990). Long-term, short-term and renegotiation: On the value of commitment in contracting. *Econometrica 58*(3), 597–619.

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

Rubinstein, A. and M. E. Yaari (1983). Repeated insurance contracts and moral hazard. *Journal of Economic Theory 30*(1), 74–97.

Sharpe, S. A. (1990). Asymmetric information, bank lending, and implicit contracts: A stylized model of customer relationships. *The Journal of Finance 45*(4), 1069–1087.

Smith, H. L. (1995). Monotone dynamical systems: an introduction to the theory of competitive and cooperative systems.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics 87*(3), 355–374.

Swinkels, J. M. (1999). Education signalling with preemptive offers. *The Review of Economic Studies 66*(4), 949–970.

Vincent, D. R. (1990). Dynamic auctions. *The Review of Economic Studies 57*(1), 49–61.

Wilson, C. (1977). A model of insurance markets with incomplete information. *Journal of Economic Theory 16*(2), 167–207.

# Appendix: Proofs

For ease of exposition, we consider (unless otherwise specified) the case of positive type monotonicity $u_{1,3} > 0$, and negative belief monotonicity $u_2 < 0$. The results are extended naturally to all other cases.

## A.1 Proof of Lemma 2

Fix an equilibrium and suppose, toward a contradiction, that there exists a pooling set: a non-degenerate interval $[\theta', \theta'']$ such that all types in this set choose the same contract $\bar{x}$. By Lemma 1, the associated prices $\bar{p}_t$ satisfy

$$\mathbb{E}[\pi(\bar{p}_t, x_t, \theta) | \theta \in [\theta', \theta'']] = 0 \quad \text{for all } t \leq T.$$

Since the zero-profit pricing function $\mathsf{p}$ is strictly increasing in $\theta$, there exists $\delta > 0$ such that

$$\bar{p}_t > \mathsf{p}(\theta'', x_t) + \delta \quad \text{for all } t \leq T.$$

We construct a deviation that strictly raises utility for type $\theta'$, assuming he has marginal utility from increasing quantities at the pooling contract. Consider $x^\varepsilon = (\bar{x}_0 - \varepsilon, \bar{x}_1, \ldots, \bar{x}_{T-1}, \bar{x}_T)$ for small $\varepsilon > 0$. If this deviation lies on-path, it must be priced as if offered by some type $\tilde{\theta} < \theta'$. To see why, suppose instead that $\tilde{\theta} > \theta''$, then type $\tilde{\theta}$ could profitably deviate to the pooling contract $\bar{x}$, obtaining simultaneously higher quantities and lower prices, contradicting equilibrium optimality. If the deviation lies off-path, then under D1 the support of dealers' beliefs must be contained in $[0, \theta')$, because $\partial \left( \sum_{t=0}^T \beta^t g(\bar{x}_t) \right) / \partial x_0 > 0$ implies lower types have strictly stronger incentives to marginally reduce $x_0$. Letting $\varepsilon \to 0$, the deviation guarantees an extra gain of at least

$$\sum_{t=0}^{T} \beta^t \left( f(\bar{p}_t - \delta, \bar{x}_t) + \theta' g(\bar{x}_t) \right) - \sum_{t=0}^{T} \beta^t \left( f(\bar{p}_t, \bar{x}_t) + \theta' g(\bar{x}_t) \right) > 0,$$

contradicting the optimality of the pooling contract for type $\theta'$. (If type $\theta'$ had negative marginal utility from increasing quantities under $\bar{x}$, the same argument applies to $\theta''$ with the sign of the deviation reversed.) □

## A.2 Proof of Lemma 3

If $\bar{x}$ is on-path, dealers perceive it as offered by the unique type $\theta$ for which $x^c(\theta) = \bar{x}$, then $\Sigma_T(\bar{x}) = \Gamma_T(\theta)$. Now suppose $\bar{x}$ is off-path. First, consider the case $\Sigma_T(\bar{x}) \in [\Gamma_T(0), \Gamma_T(1)]$. Let $\hat{\theta} \coloneqq \Gamma^{-1}(\Sigma_T(\bar{x}))$, which is well-defined because $\Gamma_T$ is strictly increasing in a separating equilibrium. We claim no other type $\theta' \neq \hat{\theta}$ can belong to the support of dealers' beliefs after $\bar{x}$. Fix such a $\theta' \neq \hat{\theta}$. There exists a price vector $\bar{p}$ such that $U(\theta') = \sum_{t=0}^T \beta^t w(\theta', \bar{p}_t, \bar{x}_t)$. Define $\varphi(\theta) \coloneqq U(\theta) - \sum_{t=0}^T \beta^t w(\theta, \bar{p}_t, \bar{x}_t)$. Since $U$ is convex and $w$ is affine in $\theta$, the function $\varphi$ is convex. Note that $\varphi(\theta') = 0$ and, using the envelope condition, $\varphi'(\theta') = \Gamma_T(\theta') - \Sigma_T(\bar{x}) \neq 0$. Hence, there exists $\theta''$ sufficiently close to $\theta'$ with $\varphi(\theta'') < 0$, that is, $U(\theta'') < \sum_{t=0}^T \beta^t w(\theta'', \bar{p}_t, \bar{x}_t)$. Thus $\theta''$ strictly prefers deviating to $\bar{x}$. By D1, whenever $\theta'$ has a weak incentive to deviate to $\bar{x}$, the presence of a nearby $\theta''$ with a strictly higher deviation gain excludes $\theta'$ from the belief support. So only $\hat{\theta}$ can remain in the support. If $\Sigma_T(\bar{x}) < \Gamma_T(0)$, then $\varphi'(\theta) = \Gamma_T(\theta) - \Sigma_T(\bar{x}) > 0$ for all $\theta \in [0, 1]$, so the deviation gain is maximized at $\theta = 0$, and beliefs must concentrate on type 0. Analogously, beliefs must concentrate on type 1 when $\Sigma_T(\bar{x}) > \Gamma_T(1)$.

## A.3 Proof of Lemma 4

On-path beliefs are pinned down by the identification map $\Gamma_T(\theta)$. Thus, each type $\theta$ chooses a contract to maximize utility subject to the identification constraint:

$$\max_{x\in\Omega} \sum_{t=0}^{T} \beta^t u(\theta,\theta,x_t) \quad \text{subject to} \quad \sum_{t=0}^{T} \beta^t g(x_t) = \Gamma_T(\theta).$$

Let $\lambda$ be the Lagrange multiplier on the constraint. The first-order condition for each $x_t$ is $u_3(\theta,\theta,x_t) = \lambda g'(x_t)$, which is independent of $t$. By Assumption 1.C, the objective is strictly concave in each $x_t$, so the solution is unique. Therefore, $x_t^c(\theta) = x_s^c(\theta)$ for all $t, s \leq T$. □

## A.4 Proof of Proposition 1

We have already established separability, so the equilibrium construction and uniqueness follow from Mailath (1987), with some important adaptations due to the multi-dimensionality of the allocation.

Incentive compatibility requires

$$\theta \in \arg\max_{\hat\theta} \sum_{t=0}^{T} \beta^t u\left(\theta,\hat\theta, x_t\left(\hat\theta\right)\right),$$

which is equivalent to the first-order condition

$$\sum_{t=0}^{T} \beta^t \left[u_2(\theta,\theta,x_t(\theta)) + u_3(\theta,\theta,x_t(\theta))\, x_t'(\theta)\right] = 0.$$

Since equilibrium quantities are constant across time (by Lemma 4), the first-order condition simplifies to the differential equation (5.2). Because the equilibrium is separating, the highest type (i.e., $\theta = 1$) must be correctly identified and cannot be credibly punished for deviating to his first-best choice, thus $x_t^c(1) = x^{FB}(1)$ (*no distortion at the top*). The initial value problem is posed at the boundary $\theta = 1$ and solved backward on $[0,1]$.

At $x^{FB}(1)$ we have $u_3(1,1,x^{FB}(1)) = 0$, so $|x'(\theta)| \to \infty$ as $x \to x^{FB}(1)$. Under Assumption 2.A (bounded $|u_2|$), the inverse differential equation is Lipschitz and thus admits a unique solution despite the singularity at the initial point. The inverse problem admits both increasing and decreasing branches, which correspond to two distinct solutions to the original problem. We claim that the only incentive-compatible solution is the one for which $x_t(\theta)$ is increasing. By convexity of $U$ and the envelope condition, $\Gamma_T(\theta)$ is non-decreasing. Differentiating gives $\Gamma_T'(\theta) = \sum_{t=0}^{T} \beta^t g'(x_t(\theta)) x_t'(\theta) \geq 0$. Hence, using that $g'(x_t) > 0$, we have that $x_t'(\theta) \geq 0$ almost everywhere, so the decreasing branch violates incentive compatibility and is ruled out. The increasing branch is then unique and extends over $[0,1]$ since $u_3 \neq 0$ away from the boundary and $u_2$ and $u_3$ are continuous.

Finally, off-path deviations are unprofitable under the identification rule in Lemma 3. For any deviation $\bar x$, let $\hat\theta(\bar x)$ be the belief implied by Lemma 3. Optimality for type $\hat\theta(\bar x)$ gives

$$U\left(\hat\theta(\bar x)\right) \geq \sum_{t=0}^{T} \beta^t u\left(\hat\theta(\bar x), \hat\theta(\bar x), \bar x_t\right) = \sum_{t=0}^{T} \beta^t \left[f\left(\mathsf{p}\left(\hat\theta(\bar x), \bar x_t\right), \bar x_t\right) + \hat\theta(\bar x) g(\bar x_t)\right].$$

By convexity of $U$,

$$U(\theta) \geq U\left(\hat\theta(\bar x)\right) + \left(\theta - \hat\theta(\bar x)\right) \Gamma_T\left(\hat\theta(\bar x)\right) = U\left(\hat\theta(\bar x)\right) + \left(\theta - \hat\theta(\bar x)\right) \Sigma_T(\bar x).$$

Combining the last two inequalities yields

$$U(\theta) \geq \sum_{t=0}^{T} \beta^t \left[f(\mathsf{p}(\hat\theta(\bar x), \bar x_t), \bar x_t) + \theta g(\bar x_t)\right] = \sum_{t=0}^{T} \beta^t u(\theta, \hat\theta(\bar x), \bar x_t),$$

so the deviation cannot be profitable.

REMARK: if application-specific bounds on quantities $x_t$ apply, the KKT conditions would imply the same optimality condition on the interior and a common binding condition across dates at the boundary, preserving time invariance. In particular, the commitment equilibrium to the dynamic insurance problem is the same if $x_t$ is restricted to $[0, 1]$, since Assumption 1 still holds when restricted to this interval. □

## A.5 Proof of Lemma 6

Our logic is of backward induction on the last period for which pooling is sustainable.

*There is no pooling in the final period $T$.* Suppose, to the contrary, that a nontrivial interval of types $\Theta \subseteq [0, 1]$ trades the same sequence $\bar{x}^{T-1}$ up to period $T-1$ and also chooses the same quantity in period $T$. Since dealers cannot distinguish types within $\Theta$ in the final period, the game at $T$ reduces to a static one. The logic of Lemma 2 therefore applies directly, and no pooling can occur.

If there is no pooling for the last $k$ periods, we show that there cannot be pooling for the last $k+1$ periods. For that purpose, Let $t < T$ be the highest period at which some types still pool, with pooling set $\bar{\Theta}$. The single-crossing condition implies that $\bar{\Theta}$ is an interval. Without loss of generality, assume it is compact. Consider a deviation by some $\theta \in \bar{\Theta}$ that mimics the equilibrium strategy $x_s(\theta)$ for all $s < t$, but offers $(x_t(\theta) - \varepsilon)$ in period $t$, for some small $\varepsilon > 0$. We now show that such $\varepsilon$-deviation is profitable for some pooling type. There are two cases to consider, as the deviation might be on or off-path.

*If the deviation is on-path, we can construct a profitable deviation for an extreme type in the pooling set.* Suppose the deviation is on-path for all small $\varepsilon$, and is priced at period $t$ as if offered by a type higher than $\max\{\theta \in \bar{\Theta}\}$. Then the highest pooling type $\theta''$ should be indifferent between remaining in the pooling set and mimicking a sufficiently close type $\theta''_+$. At period $t$, $\theta''_+$ enjoys a discrete price improvement for a marginally lower quantity. To offset this benefit, later periods would need to impose a discrete quantity reduction for $\theta''_+$ as compared to the strategy of $\theta''$. But then $\theta''$ could mimic $\theta''_+$ until period $T-1$ and deviate in the final period to choose a quantity $x_T$ such that $\sum_{t=0}^{T} \beta^t g(x_t) \in (\Gamma(\theta''), \Gamma(\theta''_+))$, thereby being correctly identified as $\theta''$ and strictly improving his payoff. A symmetric argument applies to the lowest pooling type $\theta'$ if the deviation is priced as coming from a type below $\min\{\theta \in \bar{\Theta}\}$. The logic of this claim is also naturally extended to say that if the deviation is off-path, dealers' beliefs cannot assign positive probabilities to types $\theta \notin \bar{\Theta}$.

*If the deviation is off-path, then:*

1. *dealers' beliefs must have support on a single type $\hat{\theta} \in \bar{\Theta}$.* Suppose instead that dealers assign positive probability to multiple types in $\bar{\Theta}$ offering the deviation. Then these types must be indistinguishable and face the same continuation game as they would on-path. But in that case, the gain from the deviation cannot be equal across types: the highest type $\theta'' = \max\{\theta \in \bar{\Theta}\}$ obtains the largest gain. Hence, multiple types cannot be equally willing to deviate, contradicting the assumption that the belief support is not degenerate.

2. *dealers' beliefs must assign full weight to the highest type $\theta'' = \max\{\theta \in \bar{\Theta}\}$.* Suppose instead the type most willing to deviate is $\theta < \theta''$. The continuation payoff gain for $\theta$ from period $t+1$ onward could not be positive: otherwise, by continuity and single-crossing, a nearby type $\theta + \nu$ (for $\nu > 0$) would obtain a strictly larger gain from deviating, violating the D1 restriction. So any type in $\Theta$ other than $\theta''$ must be assigned zero posterior probability.

3. *the deviation is profitable for the highest pooling type $\theta''$.* We argued above that the continuation payoff gain from the deviation cannot be positive. But the gain also cannot be negative, because the deviation yields informational gain: by deviating, the individual is correctly identified and faces less distortion. Therefore, when taking into account the full payoff from period $t$ to $T$, the deviation must be strictly profitable for type $\theta''$.

□

## A.6 Proof of Proposition 2

This proof is long, so we divide it into parts.

### A.6.1 First Order Condition

Fix an equilibrium $(\{x_t^*\}, \{\mu_{it}^*\}, \{p_{it}^*\}, \{a_t^*\})$. If the functions $\Gamma_t(\theta)$ are strictly monotonic for all $t$, they have well-defined inverses $\Gamma_t^{-1} : \mathcal{D}_t \to [0,1]$, where

$$\mathcal{D}_t = [\min\{\Gamma_t(0), \Gamma_t(1)\}, \max\{\Gamma_t(0), \Gamma_t(1)\}].$$

We do not know yet whether they are differentiable almost everywhere or not, it still might be the case that $\Gamma_t'(\theta) = 0$ in a set with positive measure. Along the equilibrium path, all choices lead to $\Sigma_t(\theta) \in \mathcal{D}_t$ by definition, so

$$\text{supp}(\mu_{it}^*(x_t, h^{t-1})) = \Gamma_t^{-1}(\Sigma_t(x_t, h^{t-1})).$$

Since we do not know yet whether $\Gamma_t^{-1}$ is differentiable, it is appropriate to change variables for taking the first-order conditions. We introduce the state variable

$$s_t = \Gamma_t^{-1}\left(\sum_{k=0}^{t} \beta^k g(x_k)\right)$$

and treat it as a choice variable together with $x_t$, imposing the law of motion

$$s_{t+1} = \Gamma_{t+1}^{-1}\left(\beta^t g(x_{t+1}) + \Gamma_t(s_t)\right).$$

The Lagrangean of the problem is

$$\mathcal{L} = \sum_{t=0}^{T} \beta^t \left[u(\theta, s_t, x_t) + \lambda_t \left(\Gamma_t(s_t) - \beta^t g(x_t) - \Gamma_{t-1}(s_{t-1})\right)\right]$$

taking FOC's

$$\begin{aligned} [x_t] \quad & u_3(\theta, s_t, x_t) - \lambda_t \beta^t g'(x_t) = 0 \\ [s_t] \quad & u_2(\theta, s_t, x_t) + (\lambda_t - \beta\lambda_{t+1})\, \Gamma_t'(s_t) = 0. \end{aligned} \tag{A.1}$$

The terminal condition $\lambda_{T+1} = 0$ implies

$$u_2(\theta, s_T, x_T) + \lambda_T \Gamma_T'(s_T) = 0.$$

It cannot be $\Gamma_T'(s_T) = 0$ nor $\lambda_T = 0$. If it were, then $u_2(\theta, s_T, x_T) = 0$, a contradiction for positive $x_T$. So we can isolate $\lambda_T$ and substitute it in the expression for $[s_{T-1}]$ to find

$$u_s(\theta, s_{T-1}, x_{T-1}) + \left(\lambda_{T-1} + \beta \frac{u_2(\theta, s_T, x_T)}{\Gamma_T'(s_T)}\right) \Gamma_{T-1}'(s_{T-1}) = 0.$$

Just like before, we cannot have $\Gamma_{T-1}'(s_{T-1}) = 0$. We proceed until $t = 0$ to conclude that no $\Gamma_t'(s_t)$ can be zero in equilibrium whenever on-path trades are positive, and they are indeed positive in a neighborhood of the initial point $\theta = 0$. Thus, we may rearrange (A.1) as

$$\begin{aligned} [x_t] \quad & \lambda_t = \frac{u_3(\theta, s_t, x_t)}{\beta^t g'(x_t)} \\ [s_t] \quad & \lambda_t = \beta\lambda_{t+1} - \frac{u_2(\theta, s_t, x_t)}{\Gamma_t'(s_t)} = -\sum_{k=t}^{T} \frac{\beta^{k-t} u_2(\theta, s_k, x_k)}{\Gamma_k'(s_k)} \end{aligned}$$

where the second equality in $[s_t]$ comes using forward substitution. By equating the two forms of $\lambda_t$ above and using that $s_t = \theta$ in equilibrium, we find the first-order conditions solely in terms of $x_t$:

$$[x_t] \qquad \sum_{k=t}^{T} \left( \frac{\beta^k u_2(\theta, \theta, x_k(\theta))}{\Gamma_k'(\theta)} \right) = \frac{-u_3(\theta, \theta, x_t(\theta))}{g'(x_t(\theta))}. \tag{A.2}$$

By first evaluating (A.2) at $t = T$ and then using backward substitution we find the system

$$\begin{aligned} \Gamma_t'(\theta) &= \frac{-\beta^t u_{2,t}}{(u_{3,t}/g_t') - (u_{3,t+1}/g_{t+1}')} \\ \Gamma_T'(\theta) &= \frac{-\beta^T u_{2,T}}{(u_{3,T}/g_T')}. \end{aligned} \tag{A.3}$$

Since $\Gamma_0(\theta) = g(x_0(\theta))$ is differentiable, $x_0(\theta)$ must also be. We can write $\Gamma_t(\theta) = \Gamma_{t-1}(\theta) + \beta^t g(x_t(\theta))$ and use forward substitution in (A.3) to conclude that all functions $x_t(\theta)$ must be differentiable as well.

Using

$$\Gamma_t'(\theta) = \sum_{k=0}^{t} \beta^k g'(x_k(\theta)) x_k'(\theta) = \Gamma_{t-1}'(\theta) + \beta^t x'(\theta) g_t'(\theta)$$

we get

$$\begin{aligned} x_0'(\theta) &= \frac{-u_{2,0}(\theta)}{u_{3,0}(\theta) - u_{3,1}(\theta) \left( \frac{g_0'(\theta)}{g_1'(\theta)} \right)} \\ x_t'(\theta) &= \frac{-u_{2,t}(\theta)}{u_{3,t}(\theta) - u_{3,t+1}(\theta) \left( \frac{g_t'(\theta)}{g_{t+1}'(\theta)} \right)} + \beta^{-1} \frac{u_{2,t-1}(\theta)}{u_{3,t-1}(\theta) \left( \frac{g_t'(\theta)}{g_{t-1}'(\theta)} \right) - u_{3,t}(\theta)} \quad \text{for } 0 < t < T \\ x_T'(\theta) &= \frac{-u_{2,T}(\theta)}{u_{3,T}(\theta)} + \beta^{-1} \frac{u_{2,T-1}(\theta)}{u_{3,T-1}(\theta) \left( \frac{g_T'(\theta)}{g_{T-1}'(\theta)} \right) - u_{3,T}(\theta)}. \end{aligned}$$

NOTATION. In what follows, we denote the system above by $x'(\theta) = F(x(\theta))$ with $x_t'(\theta) = F_t(x(\theta))$ for each $t$. For $y \in \mathbb{R}^{T+1}$ we denote by $\varphi(y)$ the solution of the system with initial condition $x(0) = y$ when $u_2 > 0$ or $x(1) = y$ when $u_2 < 0$, whenever such solution exists and is unique. To denote the value of $\varphi(y)$ at point $\theta \in [0, 1]$ we use $\varphi(y)(\theta)$.

### A.6.2 Boundary Condition

The initial condition is natural: in a separating equilibrium, the worst type can never be credibly punished for deviating, so he plays his first-best quantities. Denote the worst type by $\theta_0$, so $\theta_0 = 0$ when $u_2 > 0$ and $\theta_0 = 1$ when $u_2 < 0$. Also, denote the first-best vector as $\mathbf{x} = (x^{FB}, \ldots, x^{FB})$. The initial condition is simply $x(\theta_0) = \mathbf{x}$.

### A.6.3 Second Order Condition

Fix any contract $x$ played on-path by a type $\theta$. For each period $t$, we can rewrite (A.2) using the type-assignment functions $\hat{\theta}(x^k)$:

$$[x_t] \qquad \left( \sum_{k=t}^{T} \beta^k u_2(\hat{\theta}(x^k), \hat{\theta}(x^k), x_k) \frac{\partial \hat{\theta}(x^k)}{\partial x_t} \right) + \beta^t u_3(\hat{\theta}(x^t), \hat{\theta}(x^t), x_t) = 0.$$

Differentiating both sides with respect to $x_t$ yields

$$\begin{aligned} &\frac{\partial}{\partial \theta} \left( \left( \sum_{k=t}^{T} \beta^k u_2(\theta, \hat{\theta}(x^k), x_k) \frac{\partial \hat{\theta}(x^k)}{\partial x_t} \right) + \beta^t u_3(\theta, \hat{\theta}(x^t), x_t) \right) \Bigg|_{\theta = \hat{\theta}(x^k)} \cdot \left( \frac{\partial \hat{\theta}(x^k)}{\partial x_t} \right) \\ &\quad + \frac{\partial}{\partial x_t} \left( \left( \sum_{k=t}^{T} \beta^k u_2(\theta, \hat{\theta}(x^k), x_k) \frac{\partial \hat{\theta}(x^k)}{\partial x_t} \right) + \beta^t u_3(\theta, \hat{\theta}(x^t), x_t) \right) \Bigg|_{\theta = \hat{\theta}(x^k)} = 0. \end{aligned}$$

And rearranging we get

$$\frac{\partial^2 W(x,\theta)}{\partial x_t^2} = -\frac{\partial}{\partial \theta}\left(\left(\sum_{k=t}^{T}\beta^k u_2(\theta,\hat{\theta}(x^k),x_k)\frac{\partial\hat{\theta}(x^k)}{\partial x_t}\right)+\beta^t u_3(\theta,\hat{\theta}(x^t),x_t)\right)\Bigg|_{\theta=\hat{\theta}(x^k)}\cdot\left(\frac{\partial\hat{\theta}(x^k)}{\partial x_t}\right).$$

Our functional form restriction $w = f(p,x)+\theta g(x)$ implies $u_{12,k}=0$, so we find

$$\frac{\partial^2 W(x,\theta)}{\partial x_t^2} = -\beta^t u_{13}(\theta,\hat{\theta}(x^t),x_t)\Big|_{\theta=\hat{\theta}(x^k)}\cdot\left(\frac{\partial\hat{\theta}(x^k)}{\partial x_t}\right) = -\frac{\beta^t\left(g_t'\right)^2}{\Gamma_t'(\theta)}.$$

We conclude that $\partial\hat{\theta}(x^t)/\partial x_t$ must have the same sign as $g_t'$. We can substitute the on-path beliefs derivative to get

$$\frac{\partial^2 W(x,\theta)}{\partial x_t^2} = -\frac{\beta^t (g_{1,t})^2}{\Gamma_t'(\theta)}.$$

We concluded above that $\Gamma_t'(\theta)\neq 0$, so a necessary condition for optimization is that $\Gamma_t'(\theta)>0$ almost everywhere for all $t$.

Since the equilibrium is separating and all $\Gamma_t'>0$, we must have $x_t$ increasing in $\theta$ whenever $u_{1,3}$ is positive, and decreasing whenever $u_{1,3}$ is negative. Indeed, say that $u_{1,3}$ is positive, then higher types would rather trade more than lower types in first-best. But they are also perceived as higher types when they increase quantities, so optimality requires quantities to be increasing in type in equilibrium. The same logic applies when $u_{1,3}$ is negative. Also, this echoes the logic of Theorem 2 of Mailath (1987).

### A.6.4 Singularity

See that $F(x)$ is undefined at the initial condition $x(\theta_0)=\mathbf{x}$. It is, however, well defined in the set

$$S \coloneqq \{x\in\mathbb{R}^{T+1} : 0<x_0<\cdots<x_T<x^{FB}(0)\}.$$

In particular, the problem $x'=F(x)$ with $x(\theta_0)=y\in S$ has a unique solution, the flow $\varphi(y)$. We next present a sequence of arguments to assure existence and uniqueness of an equilibrium. First, we show that our system is monotone and positively invariant in $S$. Second, we invert our system and find a unique continuous extension of our system at the initial condition along the equilibrium path. Using such continuous extension and the positive invariance property, we show that there exists an admissible solution to the inverted system, that must also solve our original system. After that, we rely on monotonicity to prove that there is only one admissible solution to our equilibrium.

### A.6.5 Monotonicity

We need some additional definitions. First, we write $a\geq b$ if $a-b$ belongs to the nonnegative cone $\mathbb{R}_+^{T+1}$. Second, we say the system $x'=F(x)$ is *cooperative* if all off-diagonals of the Jacobian $DF$ are non-negative; and we say it is *competitive* if all off-diagonals of $DF$ are non-positive. Finally, we say the flow $\varphi(y)$ is *monotone* if the following condition holds. If $a\leq b, \theta\in(0,1)$, and $\varphi(a)(\theta)$ and $\varphi(b)(\theta)$ are defined, then $\varphi(a)(\theta)\leq\varphi(b)(\theta)$.

A direct adaptation of Proposition 1.1 in Chapter 3 of Smith (1995) to our setting says that: if $u_2>0$, a cooperative system generates a monotone dynamical system; and if $u_2<0$, a competitive system generates a monotone dynamical system. This result is known as the Kamke condition.

Thus, our system is monotone whenever all partials $\partial F_t(x)/\partial x_k$ for $k\neq 0$ have the same sign as $u_2$. Let us calculate them. First see that $\partial F_t(x)/\partial x_k = 0$ whenever $|k-t|>1$. Whenever $\partial F_t/\partial x_{t+1}$

and $\partial F_t/\partial x_{T-1}$ are well defined, we have

$$\frac{\partial F_t}{\partial x_{t+1}} = u_{2,t} \frac{\left(-g_t' \frac{\partial}{\partial x_{t+1}} \left(\frac{u_{3,t+1}}{g_{t+1}'}\right)\right)}{\left(u_{3,t}(\theta) - u_{3,t+1}(\theta)\left(\frac{g_t'(\theta)}{g_{t+1}'(\theta)}\right)\right)^2}$$

$$\frac{\partial F_t}{\partial x_{t-1}} = \frac{1}{u_{2,t-1}} \frac{u_{2,3,t-1}}{u_{1,3,t-1}} \beta^{-1}\left(-\Gamma_{t-1}'(\theta)\right) + u_{2,t-1}\beta^{-1} \frac{\left(-g_t' \frac{\partial}{\partial x_{t-1}} \left(\frac{u_{3,t-1}}{g_{t-1}'}\right)\right)}{\left(u_{3,t-1}(\theta)\left(\frac{g_t'(\theta)}{g_{t-1}'(\theta)}\right) - u_{3,t}(\theta)\right)^2}.$$

By assumption 2, both terms have the same sign as $u_2$, and our system is monotone.

For any $x \in S$, it either holds that $F_t(x) < 0$ or $F_t(x) > 0$ for all $t$. So Proposition 2.1 in Chapter 3 of Smith (1995) guarantees that $S$ is positively invariant.

### A.6.6 Unique Continuous Extension

Let us call $x_0 = s$ and define the state variables $z = (\theta, x_1, \ldots, x_T)$. The inverted system becomes

$$H_\theta(s,z) = \theta'(s) = \frac{u_{3,1}(g_0'/g_1') - u_{3,0}}{u_{2,0}}$$

$$H_t(s,z) = x_t'(s) = \frac{u_{2,t}}{u_{2,0}}\left(\frac{u_{3,1}(g_0'/g_1') - u_{3,0}}{u_{3,t+1}(g_t'/g_{t+1}') - u_{3,t}}\right) - \sum_{k=1}^{t-1} \beta^{k-t}(g_k'/g_t')x_k'(s)$$

$$H_T(s,z) = x_T'(s) = \frac{u_{2,T}}{u_{2,0}}\left(\frac{u_{3,1}(g_0'/g_1') - u_{3,0}}{-u_{3,T}}\right) - \sum_{k=1}^{T-1} \beta^{k-T}(g_k'/g_T')x_k'(s).$$

Denote the initial value $x_0(\theta_0) = s_0$. We have $\theta'(s_0) = 0$, but all the other equations have indeterminate forms of type $0/0$, because of the terms in parentheses. We apply the l'Hôspital's rule to get their limits

$$\lim_{s\to s_0} H_t(s,z) = \frac{u_{2,t}}{u_{2,0}}\left(\frac{u_{3,1}'(s_0)(g_0'/g_1') - u_{3,0}'(s_0)}{u_{3,t+1}'(s_0)(g_0'/g_1') - u_{3,t}'(s_0)}\right) - \sum_{k=1}^{t-1} \beta^{k-t} \lim_{s\to s_0}(g_k'/g_t')x_k'(s)$$

$$\lim_{s\to s_0} H_T(s,z) = \frac{u_{2,T}}{u_{2,0}}\left(\frac{u_{3,1}'(s_0)(g_0'/g_1') - u_{3,0}'(s_0)}{u_{3,T}'(s_0)}\right) - \sum_{k=1}^{T-1} \beta^{k-T} \lim_{s\to s_0}(g_k'/g_T')x_k'(s).$$

Use that $x_1(s_0) = \cdots = x_T(s_0)$ to simplify the expressions to

$$\lim_{s\to s_0} H_t(s,z) = x_t'(s_0) = \frac{x_1'(s_0) - 1}{x_{t+1}'(s_0) - x_t'(s_0)} - \sum_{k=1}^{t-1} \beta^{k-t} x_k'(s_0)$$

$$\lim_{s\to s_0} H_T(s,z) = x_T'(s_0) = \frac{x_1'(s_0) - 1}{-x_T'(s_0)} - \sum_{k=1}^{T-1} \beta^{k-T} x_k'(s_0).$$

If such limits exist, then

$$\sum_{k=0}^{t} \beta^k x_k'(s_0) = \beta^t \frac{x_1'(s_0) - 1}{x_{t+1}'(s_0) - x_t'(s_0)}$$

$$\sum_{k=1}^{T} \beta^k x_k'(s_0) = \beta^T \frac{x_1'(s_0) - 1}{-x_T'(s_0)}.$$

We have a system of $T$ variables $(x_1'(s_0), \ldots, x_T'(s_0))$ and $T$ equations. I employ a forward recursion method to show this system has a unique solution. Begin by fixing a value $\lambda = x_1'(s_0)$. Consider the pair

$$(y_1(\lambda), Y_1(\lambda)) = (\lambda, 1 + \beta\lambda).$$

Define the operator $\Psi_t(y, Y)$ as

$$\Psi_t(y, Y) \mapsto \left( y + \beta^t \frac{(\lambda - 1)}{Y}, Y + \beta^{T+1} \left[ y + \beta^t \frac{(\lambda - 1)}{Y} \right] \right),$$

then compute

$$(y_T(\lambda), Y_T(\lambda)) = \Psi_{T-1} \circ \cdots \circ \Psi_1(y_1(\lambda), Y_1(\lambda)),$$

and find $\lambda^*$ as a root of

$$K(\lambda) = Y_T(\lambda) y_T(\lambda) + \beta^T(\lambda - 1) = 0.$$

The derivatives we are looking for are found as $(x_1'(s_0), \ldots, x_T'(s_0)) = (y_1(\lambda^*), \ldots, y_T(\lambda^*))$. It is left to show that there is a unique $\lambda^*$ that solves the system. We argued before that incentive compatibility requires all functions $x_t'(\theta)$ to have the same sign. The relevant domain for our root is, therefore, $\lambda \in (0, 1)$, because $\lambda \geq 1$ would lead to $y_{t+1} \geq y_t$ for all $t$, and then $Y_T y_T + \beta^T(\lambda - 1) > 0$ for sure. Now I prove by induction that $y_t'(\lambda)$ and $Y_t'(\lambda)$ are always positive. The base case is direct: $y_1'(\lambda) = 1$ and $Y_1'(\lambda) = \beta$. For the inductive step, assume $y_t', Y_t' > 0$. Since $\lambda < 1$, we have

$$y_{t+1}'(\lambda) = y_t'(\lambda) + \beta^t \left[ \frac{Y_t(\lambda) - (\lambda - 1) Y_t'(\lambda)}{Y_t(\lambda)^2} \right] \geq y_t'(\lambda) > 0.$$

Also, since $Y_{t+1}'(\lambda) = Y_t'(\lambda) + \beta^{t+1} y_{t+1}'(\lambda)$, we have $Y_{t+1}'(\lambda) > Y_t'(\lambda) > 0$, and this concludes our induction step. It is easy to see that $K$ is continuous and strictly increasing in $\lambda \in (0, 1)$. We argued above that $K(1) > 0$, now note that $\lim_{\lambda \to 0^+} y_2(\lambda) = -\beta$, which would lead $Y_T y_T$ to be non-positive, and thus $K(0^+) < 0 < K(1^-)$, so the intermediate value theorem implies there is only one $\lambda^* \in (0, 1)$ for which $K(\lambda^*) = 0$. Once we find $\lambda_0$, we can redefine our inverted system at $s = s_0$ by its unique continuous extension: $\theta'(s_0) = 0$ and $x_t'(s_0) = y_t(\lambda^*)$ for all $t \leq T$.

Hence we can say that there exists $L \in \mathbb{R}^{T+1}$ such that any equilibrium solution starting at the specified boundary values satisfies

$$\lim_{s \to s_0} H(z(s)) = L,$$

where all entries of $L$ (except for the first-one) are nonzero and have the same sign, and $|L_{t+1}| < |L_t|$ for all $t < T$.

### A.6.7 Existence

Let us work with the inverted system where the time-variable is $s = x_0$. Suppose, without loss, that we are interested in solutions with $s$ ranging from $s_0 = x^{FB} > 0$ to 0. The system is

$$z'(s) = H(z(s)) \text{ for }, z'(s_0) = L, \text{ and } z(s_0) = \bar{z} = (\theta_0, x^{FB}, \ldots, x^{FB}),$$

where $L \coloneqq (0, y_0(\lambda^*), \ldots, y_T(\lambda^*))$ is the unique continuous extension along the equilibrium path that we found above.

We know at $s_0$ the system is at $\bar{S}$, but $L$ points strictly towards the open set $S$. So any solution of the extended inverted system enters $S$ immediately after $s_0$ and, by the positive invariance, stays in $S$ thereafter. Define, for small $\Delta > 0$, the functions $z_\Delta$ where $z_\Delta(s) = \bar{z} + (s - s_0)L$ for $s \in [s_0 - \Delta, s_0]$ and $z_\Delta(s)$ is the unique solution to $z' = H(z)$ with $z(s_0 - \Delta) = \bar{z} + (\Delta - s_0)L$ at $s \in [0, s_0 - \Delta]$. Then take $\Delta \to 0$. The solutions are extended back to the initial point $\bar{z}$ by continuity, thus satisfying $z' = H(z)$ on the whole interval $(0, s_0)$, and their limiting derivatives at $s_0$ coincide with $L$ by construction. Inverting the system back we have a solution for $x' = F(x)$ with $x(\theta_0) = \mathbf{x}$.

### A.6.8 Uniqueness

Consider the case $u_2 > 0$, so initial condition is $x(0) = \mathbf{x}$. By the previous argument involving the second-order condition, admissible solutions involve $x_t'(\theta) < 0$ for all $t$, which is consistent with $x(\theta) \in S$ for all $\theta \in (0, 1)$. The case $u_2 < 0$ follows analogously. We proceed in two claims.

*Claim 2.1:* for any admissible solution $\tilde{x}$ there exists a monotone sequence $\tilde{y}_n \to \mathbf{x}$ with $\tilde{y}_n \in S$ for every $n$, and $\varphi(\tilde{y}_n) \to \tilde{x}$ uniformly on every $[0, 1-\varepsilon], \varepsilon > 0$.
*Proof:* Pick a strictly decreasing sequence $\tau_n \downarrow 0$ with $\tau_n \in (0,1)$ for every $n$ and $\tau_0 \leq \varepsilon$. Set $\tilde{y}_n \coloneqq \tilde{x}(\tau_n) \in S$. Since the system is autonomous,

$$\varphi(\tilde{y}_n)(\theta) = \tilde{x}(\theta + \tau_n) \quad \forall \theta \in [0, 1-\tau_n].$$

We know $\tilde{x}$ is continuous. Since $x_t(\theta)$ is non-increasing for every $t$, we have $\tilde{y}_{n+1} \geq \tilde{y}_n$ and $\varphi(\tilde{y}_n)(\theta)$ increases pointwise to $\tilde{x}(\theta)$. Dini's theorem states that if $\varphi(\tilde{y}_n)$ converges monotonically to a continuous function $\tilde{x}$ in a compact $[0, 1-\tau_0]$, then it converges uniformly to $\tilde{x}$. □

*Claim 2.2*: If $\tilde{x}$ and $\hat{x}$ are admissible, then $\tilde{x} = \hat{x}$ on $[0,1]$.
*Proof:* let $\tilde{y}_n$ and $\hat{y}_n$ be increasing sequences on $S$ with limit $\mathbf{x}$ such that $\varphi(\tilde{y}_n) \to \tilde{x}$ and $\varphi(\hat{y}_n) \to \hat{x}$ uniformly. Choose a subsequence $z_n$ of $\tilde{y}_n$ with $z_n \geq \hat{y}_n$ for all $n$ (which is possible by Auxiliary Claim below). By order preservation, $\varphi(z_n) \geq \varphi(\hat{y}_n)$. Passing to limits along $[0, 1-\varepsilon]$ gives $\tilde{x} \geq \hat{x}$. Swap roles to get $\hat{x} \geq \tilde{x}$. We must have equality. □

*Auxiliary Claim.* Let $O$ be a nonempty and bounded open set in $\mathbb{R}^n$, and $\mathbf{y} \in \bar{O} \setminus O$ with $\mathbf{y} \gg y$ for all $y \in O$. For any two increasing sequences $\tilde{y}_n$ and $\hat{y}_n$ in $O$ with $\tilde{y}_n, \hat{y}_n \to \mathbf{y}$, there is a subsequence $z_n$ of $\tilde{y}_n$ such that $z_n \geq \hat{y}_n$ for all $n$.
*Proof:* Let $\varepsilon_1 \coloneqq ||\mathbf{y} - \hat{y}_1||_\infty$. Take $n_1$ such that $N > n_1$ ensures $||\mathbf{y} - \tilde{y}_N|| \leq \varepsilon$, and set $z_1 = \tilde{y}_{N_1}$ for some $N_1 > n_1$. Having defined $\{z_n\}_{n=1}^{k-1}$, set $\varepsilon_k \coloneqq ||\mathbf{y} - \hat{y}_k||_\infty$. There is $n_k$ such that $N_k > n_k$ such that $||\mathbf{y} - \tilde{y}_{N_k}|| \leq \varepsilon_k$. Set $N_k > \max\{n_k, N_{k-1}\}$. □

## A.7 Auxiliary Result to Proposition 3

AUXILIARY LEMMA 1. *If the* $\Gamma$*–rule holds for all* $t < k \leq T$*, then the unique solution to* (6.4) *is the unique* $\theta$ *such that* $\Gamma_t(\theta) = \Sigma_t$.

*Proof.* Denote the right-hand side of (6.4) by $\tilde{\Sigma}_T\left(\hat{\theta} \,\middle|\, x_t, h^{t-1}\right)$, interpreted as the optimal terminal signature for type $\hat{\theta}$ following $(x_t, h^{t-1})$.

We proceed by backward induction.

Base case: $\tilde{\Sigma}_T = \Gamma_T(\theta)$ is optimal in the final period following $h^{T-1}$ for type $\theta$ only if $\Sigma_{T-1} = \Gamma_{T-1}(\theta)$. Indeed, suppose $\Gamma_{T-1}(\theta) \neq \Sigma_{T-1}$, and the optimal choice satisfies $\Sigma_T = \Gamma_T(\theta)$. Since $\Sigma_T = \Sigma_{T-1} + \beta^T g(\tilde{x}_T)$, and $g$ is strictly monotone, this pins down a unique

$$\tilde{x}_T = g^{-1}\left(\frac{\Gamma_T(\theta) - \Sigma_{T-1}}{\beta^T}\right).$$

Because $\Gamma_T(\theta) = \Gamma_{T-1}(\theta) + \beta^T g(x_T^*(\theta))$, whenever $\Sigma_{T-1} \neq \Gamma_{T-1}(\theta)$ we have $\tilde{x}_T \neq x_T^*(\theta)$. Hence the first-order condition

$$\frac{\beta^T u_2(\theta, \theta, \tilde{x}_T) g'(\tilde{x}_T)}{\Gamma_T'(\theta)} + u_3(\theta, \theta, \tilde{x}_T) = 0 \tag{A.4}$$

would have to hold at a point different from $x_T^*(\theta)$. However, under equilibrium optimality, this equation should have a unique solution at $x_T^*(\theta)$, which is a contradiction.

Inductive hypothesis: $\tilde{\Sigma}_T = \Gamma_T(\theta)$ is optimal following $h^k$ for type $\theta$ only if $\Sigma_k = \Gamma_k(\theta)$ for all $k \in \{t, \ldots, T\}$. We now show that $\tilde{\Sigma}_T = \Gamma_T(\theta)$ is optimal following $h^{t-1}$ for type $\theta$ only if $\Sigma_{t-1} = \Gamma_{t-1}(\theta)$. See that if $\tilde{\Sigma}_T = \Gamma_T(\theta)$ is reached, then two conditions must be met. First, $\Sigma_t = \Gamma_t(\theta)$ must be reached (by the inductive hypothesis), so $\beta^t \tilde{x}_t = \Gamma_t(\theta) - \Sigma_{t-1}$. Second, $\tilde{x}_t$ must satisfy the first-order condition at $t$:

$$\sum_{k=t}^{T} \frac{\beta^{k-t} u_2(\theta, \theta, \tilde{x}_k) g'(\tilde{x}_k)}{\Gamma_k'(\theta)} + u_3(\theta, \theta, \tilde{x}_t) = 0. \tag{A.5}$$

Note that by the inductive hypothesis, all $\tilde{x}_k$ from $t+1$ until $T$ are fixed and must coincide with $x_k^*(\theta)$. The two conditions can only be true together if $\Sigma_{t-1} = \Gamma_{t-1}(\theta)$. This completes the proof of the claim.

See that equations (A.4) and (A.5) have a unique solution because the separating equilibrium that satisfies the $\Gamma$ equilibrium is uniquely determined and features $\Gamma_t' > 0$ for all $t$. We are actually extending the $\Gamma$ equilibrium in our proposition to an equilibrium that satisfies the $\Gamma$–rule for periods $k \geq t$, but it's constructed just like in Lemma 7. □

## A.8 Proof of Lemma 7

Consider the insurance setting. Denote by $U^c(\theta)$ and $U^*(\theta)$ the value functions for type $\theta$ under commitment and no-commitment, respectively. Both are continuous, increasing and convex, and $U^c(1) = U^*(1)$. Abusing notation, denote by $U^c_\theta(x)$ and $U^*_\theta(x)$ the value functions of type $\theta$ from playing $x$, assuming the $\Gamma$-criterion holds in each respective equilibrium.

First, there is $\varepsilon > 0$ such that $U^c(\theta) \geq U^*(\theta)$ for all $\theta \in (1-\varepsilon, 1]$. That is, types close to 1 are better off under the commitment equilibrium. If this was not the case, there would be $\hat{\varepsilon} > 0$ such that $x^*(\theta) \succ_\theta x^c(\theta)$ for all $\theta \in (1-\hat{\varepsilon}, 1)$ under fair pricing. But then all these types could replicate these allocations under commitment and still obtain separation and fair prices, thus contradicting the optimality assumption we began with for the commitment equilibrium.

The other relevant case is when there is $\hat{\theta} \in (0, 1]$ where $U^c(\theta) - U^*(\theta)$ changes sign (from positive on $\theta > \hat{\theta}$ to negative on $\theta < \hat{\theta}$). Since $U^c(\theta) - U^*(\theta)$ is differentiable, the derivatives of the value function at $\hat{\theta}$ must be distinct. In particular,

$$\frac{\partial U^c(\hat{\theta})}{\partial \theta} > \frac{\partial U^*(\hat{\theta})}{\partial \theta} \iff \Gamma^c_T(\hat{\theta}) > \Gamma^*_T(\hat{\theta}).$$

But the $\Gamma$-rule and the convexity of the value functions then imply that playing $x^*(\hat{\theta})$ under commitment yields a strictly better pricing than playing $x^c(\hat{\theta})$. And we conclude that

$$U^c(x^*(\hat{\theta})) \geq U^c(x^c(\hat{\theta})) = U^*(x^*(\hat{\theta})),$$

which contradicts optimality again.

## A.9 Proof of Lemma 8

In the dynamic insurance case, the system in Proposition 2 simplifies to

$$\begin{aligned}
x_t'(\theta) &= \frac{x_t(\theta)}{\gamma\sigma^2(x_{t+1}(\theta) - x_t(\theta))} - \sum_{k=0}^{t-1} \beta^{k-t}\, x_k'(\theta), \quad \text{for all } t < T \\
x_T'(\theta) &= \frac{x_T(\theta)}{\gamma\sigma^2(1 - x_T(\theta))} - \sum_{k=0}^{T-1} \beta^{k-T} x_k'(\theta) \\
x_t(1) &= 1, \quad \text{for all } t < T.
\end{aligned}$$

For comparative statics, we call $\nu \coloneqq \gamma\sigma^2$ and $\delta \coloneqq \beta^{-1}$ and incorporate such parameters as initial conditions. The system in the form $x' = f(x)$ with $x = (x_0, \ldots, x_T, \nu, \delta) \in \mathbb{R}^{T+3}$ is

$$\begin{aligned}
x_0' = f_0 &= \tfrac{1}{\nu} \frac{x_0}{x_1 - x_0} \\
&\vdots \\
x_t' = f_t &= \tfrac{1}{\nu} \left( \frac{x_t}{x_{t+1} - x_t} - \delta \frac{x_{t-1}}{x_t - x_{t-1}} \right) \\
&\vdots \\
x_T' = f_T &= \tfrac{1}{\nu} \left( \frac{x_T}{1 - x_T} - \delta \frac{x_{T-1}}{x_T - x_{T-1}} \right) \\
\nu' &= f_{T+1} = 0 \\
\delta' &= f_{T+2} = 0.
\end{aligned}$$

We claim that our system is competitive. For that, it suffices to see that the opposite of the Jacobian $Df$ is Metzler, i.e., all off-diagonal components of $Df$ are non-positive:

$$Df = \begin{bmatrix} \nabla f_0 & \dots & \nabla f_{T+2} \end{bmatrix}^T$$
$$= \begin{bmatrix} \frac{\partial f_0}{\partial x_0} & \frac{-x_0}{\nu(x_1-x_0)^2} & 0 & 0 & \dots & 0 & 0 & \frac{\partial f_0}{\partial \nu} & \frac{\partial f_0}{\partial \delta} \\ \frac{-\delta x_1}{\nu(x_1-x_0)^2} & \frac{\partial f_1}{\partial x_1} & \frac{-x_1}{\nu(x_2-x_1)^2} & 0 & \dots & 0 & 0 & \frac{\partial f_1}{\partial \nu} & \frac{\partial f_1}{\partial \delta} \\ 0 & \frac{-\delta x_2}{\nu(x_2-x_0)^2} & \frac{\partial f_2}{\partial x_2} & \frac{-x_2}{\nu(x_3-x_2)^2} & \dots & 0 & 0 & \frac{\partial f_2}{\partial \nu} & \frac{\partial f_2}{\partial \delta} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \frac{-\delta x_T}{\nu(x_T-x_{T-1})2} & \frac{\partial f_T}{\partial x_T} & \frac{\partial f_T}{\partial \nu} & \frac{\partial f_T}{\partial \delta} \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Since $\frac{\partial f_i}{\partial \nu}, \frac{\partial f_i}{\partial \delta} < 0$ for all $i \leq T$, the system is competitive. We are interested in the time-reversed flow of the system, since our boundary condition is at $\theta = 1$. We know that the time-reversed flow of a competitive system is monotone. This is Proposition 1.1 in Chapter 3 of Smith (1995), which we already discussed in the proof of Proposition 2. The Lemma is an immediate corollary of this proposition.

REMARK 1. Although Picard–Lindelöf does not apply to our system, we know there is a unique admissible solution found constructively in the proof of Proposition 2. So build a sequence of uniquely defined flows converging to such solution and apply the monotone comparison term by term (changing only $\nu$ or $\delta$). The monotonicity must also hold at the limiting function, so our results are valid.

REMARK 2. For the general case, without specifying functional forms, we would need a prior step of checking monotonicity of all the initial conditions with respect to each parameter. This step is not needed for our insurance specification because for any $(\beta, \gamma, \sigma^2)$ the initial condition is full-insurance for the riskiest type.

## A.10 Formal Result for the Extension

LEMMA 9. $0 < \bar{\theta}^c < \bar{\theta}^*$.

*Proof.* In the first phase:

$$x_0'(\theta) = \frac{-v'(\theta)(\phi(x_0(\theta)) + \beta\phi(1 - x_0(\theta)))}{(v(\theta)\phi'(x_0(\theta)) - \theta) - \beta(v(\theta)\phi'(1 - x_0(\theta)) - \theta)}$$

and $\Gamma_1(\theta) = x_0(\theta)(1 + \beta)$. The payoff for type $\theta$ from trading $x_1(\theta) - \varepsilon$ at $t = 1$ for a small $\varepsilon > 0$ is

$$\xi_\theta(\varepsilon) \coloneqq v\left(\Gamma_1^{-1}(x_0(\theta)(1-\beta) + \beta(1-\varepsilon))\right)\phi(1 - x_0(\theta) - \varepsilon) - \theta(1 - x_0(\theta) - \varepsilon)$$

$$\xi_\theta'(0) = -[\phi'(x_1(\theta))v(\theta) - \theta] - \frac{\beta\phi(x_1(\theta))v'(\theta)}{\Gamma_1'(\theta)}.$$

Substituting the expression for $\Gamma_1'(\theta)$, we have

$$\xi_\theta'(0) = -[\phi'(x_1(\theta))v(\theta) - \theta] + \frac{\beta}{1-\beta}\left(\frac{\phi(x_1(\theta))}{\phi(x_0(\theta)) + \beta\phi(x_1(\theta))}\right)[\phi'(x_0(\theta))v(\theta) - \theta) - \beta(\phi'(x_1(\theta))v(\theta) - \theta)].$$

When $x_0(\theta) = x_1(\theta) = \frac{1}{2}$, we have

$$\xi_\theta'(0) = [\phi'(\tfrac{1}{2})v(\theta) - \theta]\left[\frac{\beta}{1+\beta} - 1\right] < 0.$$

When $x_0(\theta) \to 0$ we have

$$\xi_\theta'(0) = \frac{v(\theta)}{1-\beta}[\phi'(0) - \phi'(1)] > 0.$$

□