

SUPPLEMENT TO “LONG-TERM CONTRACTING WITH
TIME-INCONSISTENT AGENTS”

(*Econometrica*, Vol. 89, No. 2, March 2021, 793–824)

DANIEL GOTTLIEB

Department of Management, London School of Economics

XINGTAN ZHANG

Leeds School of Business, University of Colorado Boulder

APPENDIX A: CONSUMER HETEROGENEITY

IN THIS APPENDIX, we present the results summarized in Section 3.3. To simplify the exposition, we assume that the consumer has a constant deterministic income w in each period. The setting is the same as in Section 2, except that firms do not know the consumer’s “type” (either $\hat{\beta}$ or β). Instead, they have some prior distribution over the consumer’s possible types.

This is a dynamic game with incomplete information, where the consumer’s contract can signal his type to the firm. After seeing the contract offered by the consumer, the firm must update its beliefs about the consumer’s type. We therefore incorporate the standard consistency condition from perfect Bayesian equilibrium, which requires the firm’s beliefs to be consistent with Bayes’s rule on histories that are reached with positive probability.

A pure strategy for type θ of the time-1 self is a consumption vector (“contract”) $\mathbf{c}(\theta)$. A pure strategy for the firm is a mapping d from the space of possible contracts to $\{0, 1\}$, which specifies whether the firm accepts ($d = 1$) or rejects ($d = 0$) each contract offered by the time-1 self.

A time- t history describes all actions by the consumer and all uncertainty realized until period t : $\mathbf{h}_t = (\mathbf{c}, h^t, s_t)$, where $h^t \in \{A, B\}^{t-1}$ is an option history (as defined in Section 2). Let \mathcal{H}_t denote the set of all possible time- t histories. A pure strategy for type θ of self $t \in \{2, \dots, T - 1\}$ is a mapping from the history of previous actions and realized uncertainty to an action $\sigma_t : \mathcal{H}_t \rightarrow \{A, B\}$.²⁹

We can now generalize the definition of perception-perfect equilibrium (see Appendix E) to incorporate imperfect information. A “perception-perfect equilibrium with Bayesian rationality on the firm side” (henceforth *equilibrium*) is a contract and a pair of strategies for each consumer type,

$$(\mathbf{c}(\theta), \sigma_2(\theta), \dots, \sigma_{T-1}(\theta)) \quad \text{and} \quad (\hat{\sigma}_2(\theta), \dots, \hat{\sigma}_{T-1}(\theta)),$$

and a strategy d for the firm such that:

- For each θ , $\mathbf{c}(\theta)$ maximizes the expected experienced utility (E1) of the time-1 self of type θ under the assumption that each self $r > 1$ uses strategy $\hat{\sigma}_r(\theta)$ and the firm uses strategy d .

Daniel Gottlieb: d.gottlieb@lse.ac.uk

Xingtang Zhang: xingtang.zhang@colorado.edu

²⁹As in the model without heterogeneity, there is no loss of generality in assuming that there is only one option in period T . Therefore, we do not need to include choices by self T .

- For each θ , history h_t , and $t > 1$, the strategy $\sigma_t(\theta)$ maximizes the expected experienced utility (E1) of the time- t self of type θ conditional on h_t under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r(\theta)$.
- For each θ , history h_t , and $t > 1$, the strategy $\hat{\sigma}_t(\theta)$ maximizes the expected perceived utility (E2) of the time- t self of type θ conditional on h_t under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r(\theta)$.
- For each \mathbf{c} , the acceptance decision $d(\mathbf{c})$ maximizes the firm's expected discounted profits given the firm's beliefs about the consumer's type θ under the assumption that each type of the consumer uses strategy $\sigma_t(\theta)$ in all periods $t > 1$.
- For any contract offered by some type $\tilde{\mathbf{c}} \in \mathbf{c}(\Theta)$, the firm's beliefs about the consumer's type θ are derived by Bayes's rule.

A.1. *Unknown Naiveté*

In this subsection, we show that the results from Section 2 remain unchanged when the firm does not know the consumer's naiveté parameter $\hat{\beta}$. Suppose the firm has a prior distribution with full support over the non-degenerate type space $\Theta \subseteq (\beta, 1]$. Note that the type space may be discrete or continuous.

We refer to the equilibrium consumption in the model in which the firm knows the consumer's type (Section 2) as the “full-information contract.” We now show that any equilibrium of this game has complete pooling at the full-information contract:

PROPOSITION 4: *Suppose the firm does not know the consumer's naiveté parameter. There exists an equilibrium. Moreover, in any equilibrium, there is complete pooling at the full-information contract.*

The key intuition for Proposition 4 is that the equilibrium contract with full information does not depend on $\hat{\beta}$ (Corollary 1). Since that contract maximizes each type's perceived utility and gives zero profits, there are no beliefs by the firm about the consumer's type that would allow the consumer to obtain a higher perceived utility while not giving negative profits to the firm.

Proposition 4 implies that, as in Theorem 1, the welfare loss from dynamic inconsistency vanishes as the contracting horizon grows.

A.2. *Adding Sophisticated Consumers*

In the previous subsection, we assumed that, while the firm did not know the consumer's naiveté parameter $\hat{\beta}$, it still knew that the consumer was (at least partially) naive, so that $\hat{\beta} > \beta$ for all types. We now introduce a sophisticated consumer type into the analysis. Formally, consider the type space $\Theta \subseteq [\beta, 1]$ and suppose that the support of the firm's beliefs about the consumer's naiveté parameter includes both the sophisticated type ($\hat{\beta} = \beta$) and at least one naive type ($\hat{\beta} > \beta$). Note that, as before, the type space can be discrete or continuous.

The proposition below shows that in the equilibrium of this game, all naive types get the full-information contract:

PROPOSITION 5: *Suppose the support of the firm's beliefs includes both the sophisticated type and at least one naive type. Then, in any equilibrium, all types get their full-information contract.*

The intuition is as follows. When given the naive consumers' full-information contract, the sophisticated consumer understands that his future selves will pick the alternative, rather than the baseline option. So he prefers to offer his own full-information contract, which prevents his future selves from deviating. However, because naive consumers believe they will pick the baseline option, they have a higher perceived utility from their own full-information contract. And because they are each offered their full-information contracts, there are no beliefs that firms can have about consumer types that would justify them offering any other contract. Then, as in the model with full information, the inefficiency from time inconsistency vanishes for all naive consumers as the contracting horizon grows (but not for the sophisticated consumer).

A.3. Unknown Time-Consistency Parameter

Suppose now that the firm does not know the consumer's time-consistency parameter β . The firm has a prior distribution with full support over the non-degenerate type space $\Theta \subseteq (0, \hat{\beta}]$. When $\hat{\beta} \in \Theta$, the model allows for both a sophisticated time-inconsistent type ($\hat{\beta} = \beta < 1$) and a time-consistent type ($\hat{\beta} = \beta = 1$). To avoid situations in which consuming all resources in the first period maximizes welfare (which would coincide with a present-biased consumer's choice), we assume that $\lim_{c \searrow 0} u'(c) = +\infty$ in this subsection.

We first show that, unlike when the consumer's naiveté is not known, there is no equilibrium in which multiple types get their full-information contracts:

LEMMA 5: *There exists \bar{T} such that for all $T > \bar{T}$, any equilibrium has at most one type offering his full-information contract.*

The intuition of Lemma 5 is as follows. If more than one type offered their full-information contracts, the more time-consistent one would pick the full-information contract of the less time-consistent type and choose the baseline rather than the alternative option. The firm offering this contract would lose money, since the baseline was an unprofitable decoy option not meant to be chosen on the equilibrium path.

Having shown that the full-information contracts cannot be supported in equilibrium, we now turn to the characterization of the equilibrium.³⁰ As is common in signaling games, if we do not impose restrictions on beliefs that firms can have off the equilibrium path, there are many equilibria. We adopt the D1 criterion (Banks and Sobel (1987)) to deal with this multiplicity issue. For simplicity, we assume that there are only two types: β_L with probability q and β_H with probability $1 - q$, where $0 < \beta_L < \beta_H \leq \hat{\beta} \leq 1$. Note that when $\beta_H = \hat{\beta} = 1$, the model has one time-inconsistent type and one time-consistent type. It is straightforward to extend our results to any finite number of types.

We will show that the allocation in the unique equilibrium that survives D1 corresponds to the "least-costly separating allocation." In this allocation, the high type gets the full-information consumption, whereas the low type gets the allocation that maximizes his perceived utility among those leaving zero profits to the firm and ensuring that the high

³⁰Although this appendix assumes that the consumer has the bargaining power, which leads to a signaling model, a similar result can be shown when the bargaining power is on the firm side. In that case, the model becomes one of screening. It can be shown that, for generic income paths, the full-information contract is not incentive compatible as the contract length grows.

type does not wish to deviate. That is, the allocation of the high type solves program (6), whereas the allocation of the low type solves the following program:

$$\max_{\{c_{(s_t)}\}} u(c_1) + \beta_L E \left[\sum_{t=2}^T \delta^{t-1} u(c_t(B, B, \dots, B)) \right], \quad (\text{A1})$$

subject to zero-profit condition, (IC), (PC), and the constraint that requires type β_H to prefer his full-information contract over type β_L 's contract:

$$\begin{aligned} & u(c_1) + \beta_H E \left[\sum_{t=2}^T \delta^{t-1} u(c_t(B, B, \dots, B)) \right] \\ & \leq u(c_1^H) + \beta_H E \left[\sum_{t=2}^T \delta^{t-1} u(c_t^H(B, B, \dots, B)) \right]. \end{aligned} \quad (\text{A2})$$

From the previous lemma, it follows that (A2) must bind when T is large enough, so the equivalence to the auxiliary program no longer holds for the low type.

The next lemma formally shows that the high type gets his full-information contract in any equilibrium that survives D1 (i.e., there is no distortion at the top):

LEMMA 6: *In any equilibrium that survives D1, type β_H gets his full-information contract.*

We next show that there is an equilibrium that survives D1 in which consumers get the least-costly separating allocation. Furthermore, we show that the equilibrium that survives D1 is unique.

LEMMA 7: *There exists an equilibrium that survives D1 in which consumers get the least-costly separating allocation. Furthermore, in any equilibrium surviving D1, the consumers get the least costly separating allocation.*

Let W_T^L denote the equilibrium welfare of type β_L , and recall that W_T^C is the equilibrium welfare of the time-consistent consumer (who maximizes welfare). The proposition below establishes that in equilibrium, the less time-consistent type gets a contract in which he consumes more than the full-information amount in the first period, thereby under-saving for the future. Moreover, this informational distortion does not vanish as the contracting length grows, so his equilibrium allocation does not converge to the Pareto frontier:

PROPOSITION 6: *Suppose $\lim_{c \searrow 0} u'(c) = +\infty$. In any equilibrium satisfying D1:*

- *There exists \bar{T} such that, for all $T > \bar{T}$, type β_L consumes more in the first period than in his full-information contract.*
- *The welfare loss is uniformly bounded away from zero as the contracting horizon grows: $\liminf_{T \nearrow \infty} (W_T^C - W_T^L) > 0$.*

APPENDIX B: NON-EXCLUSIVE CONTRACTS

This appendix considers the model in which contracts are not exclusive, so consumers can, at any point in time, sign a new contract with another firm. As in the model with one-sided commitment, to characterize the equilibrium consumption, there is no loss of

generality in restricting attention to equilibria in which the consumer never contracts with another firm.³¹

When contracts are not exclusive, firms cannot add unprofitable baseline options that naive consumers think they will choose but end up not choosing. If they offered such a contract, the consumer would stick to the baseline and readjust consumption by contracting with another firm. Therefore, any equilibrium contract must make zero profits along both the consumer's perceived path and the equilibrium (i.e., firm's perceived) path. In fact, our next lemma shows that, starting at any history, the expected PDV of future consumption must be the same in all option histories:

LEMMA 8: *Suppose contracts are not exclusive. For any (s_t, h^t) , the expected present discounted value of consumption in any option history path following h^t must be the same.*

The proof is in the supplementary appendix, but its intuition is straightforward. With non-exclusive contracts, the consumer can always smooth consumption by contracting with a new firm. Therefore, he would always pick the option path with the highest PDV of consumption.

Consider an (auxiliary) consumption-savings problem, in which the consumer is endowed with the expected PDV of income $\{w(s_t)\}$ in period 1. The only asset available is a risk-free bond that pays a gross return R , and the consumer can freely save or borrow. Since the consumer is time-inconsistent and naive, each period's self decides how much to consume and underestimates the present bias of future selves. As before, we focus on perception-perfect equilibria of this game.

To obtain the equilibrium consumption, we need to specify both how much the agent thinks his future selves will consume and how much they actually consume. Let a_t denote the asset available to the agent at time 1: $a_1 \equiv E \sum_{t=1}^T \frac{w(s_t)}{R^{t-1}}$. The agent, who has a_t asset at time t and believes he will choose consumption in periods $s > t$ according to $\hat{c}(a_s)$, believes that in period t , he will consume

$$\hat{c}_t(a_t) \in \arg \max_{\tilde{c}} u(\tilde{c}) + \hat{\beta} \sum_{s>t} \delta^{s-t} u(\hat{c}_s(a_s)), \quad (\text{B1})$$

subject to

$$\tilde{c} + \sum_{s>t} \frac{\hat{c}_s(a_s)}{R^{s-t}} \leq a_t, \quad (\text{B2})$$

$$a_{t+1} = R(a_t - \tilde{c}), \quad (\text{B3})$$

$$a_{s+1} = R(a_s - c_s(a_s)) \quad \text{for all } s > t. \quad (\text{B4})$$

However, in period t , he chooses to consume

$$c_t(a_t) \in \arg \max_{\tilde{c}} u(\tilde{c}) + \beta \sum_{s>t} \delta^{s-t} u(\hat{c}_s(a_s)), \quad (\text{B5})$$

subject to (B2), (B3), and (B4).

³¹We assume that contracting is costless. If the cost of contracting with another firm is large enough, we return to the baseline model in which consumers can commit to long-term contracts. More generally, one can envision situations in which the cost of contracting is positive but not too large, so consumers only have partial commitment.

The next proposition establishes the equivalence between non-exclusive contracts and the consumption-savings problem.

PROPOSITION 7: *The problem with non-exclusive contracts is equivalent to the consumption-saving problem. In particular, the consumption paths in the two problems are the same.*

Last, we show that the welfare loss in the consumption-savings problem does not vanish as the contracting horizon goes to infinity. Note that if in the welfare-maximizing allocation the agent consumes all resources in the first period, leaving zero consumption in all future periods, there is no scope for contracting with other firms after the first period. Then, there is no welfare loss from non-exclusive contracting. To rule out this uninteresting case, we proceed as in Section A.3 and assume that $\lim_{c \searrow 0} u'(c) = +\infty$.

PROPOSITION 8: *Suppose contracts are not exclusive, u is bounded, $\delta < 1$, and $\lim_{c \searrow 0} u'(c) = +\infty$. The welfare loss from time inconsistency is uniformly bounded away from zero as the contracting horizon T goes to infinity.*

APPENDIX C: EFFORT

This appendix formally presents the analysis from Section 3.6, where we considered the effort model. As in the consumption model, when the agent is naive, contracts involve two options in each period: a baseline effort (B) that the agent thinks his future selves will pick and an alternative effort (A) that they end up picking. As before, let h^t denote the options chosen by the agent up to time t .

The effort path of the naive agent solves

$$\min_e \sum_{t=1}^T D_{t-1} C(e_t(B, B, \dots, B))$$

subject to

$$\sum_{t=1}^T e_t(A, \dots, A) = E_T, \quad (\text{C1})$$

$$\sum_{t=\tau}^T D_{t-\tau} C(e_t(A, B, \dots, B)) \leq \sum_{t=\tau}^T D_{t-\tau} C(e_t(B, B, \dots, B)), \quad \forall \tau \geq 2, \quad (\text{C2})$$

$$\sum_{t=\tau}^T \hat{D}_{t-\tau} C(e_t(B, B, \dots, B)) \leq \sum_{t=\tau}^T \hat{D}_{t-\tau} C(e_t(A, B, \dots, B)), \quad \forall \tau \geq 2. \quad (\text{C3})$$

That is, the agent minimizes his perceived discounted cost subject to the task-completion constraint (C1), IC (C2), and PC (C3). This program is analogous to the one in the proof of Proposition 1, except that the zero profits constraint is replaced by the task completion constraint (C1) and the agent minimizes his discounted effort costs rather than maximizes his discounted utility. As before, PC requires the agent to believe that his future selves pick B, whereas IC requires them to switch to A instead.

Since the firm and the agent disagree on the options that the agent will pick, they have different beliefs about the total effort that will be exerted on the equilibrium path. The

firm accepts a contract as long as it believes that the agent will complete the task, regardless of what the agent believes. Therefore, as with the zero profits constraint in Proposition 1, the task-completion constraint (C1) only needs to hold according to the firm's beliefs.

We solve this program in the proof of Proposition 3 in the appendix. Here we illustrate it by solving the case of three periods and quasi-hyperbolic discounting, as we did in the text for Lemma 2. This illustration helps clarify the difference between contracting over consumption and over effort.

The equilibrium program becomes

$$\min_e C(e_1) + \beta[\delta C(e_2(B)) + \delta^2 C(e_3(B))]$$

subject to

$$e_1(A) + e_2(A) + e_3(A) = E_3, \quad (\text{C4})$$

$$C(e_2(B)) + \hat{\beta}\delta C(e_3(B)) \leq C(e_2(A)) + \hat{\beta}\delta C(e_3(A)), \quad (\text{PC})$$

$$C(e_2(A)) + \beta\delta C(e_3(A)) \leq C(e_2(B)) + \beta\delta C(e_3(B)). \quad (\text{IC})$$

Note first that (IC) must bind. Otherwise, we could reduce the perceived cost in the objective function by reducing $e_3(B)$. Since (IC) binds, (PC) can be written as a monotonicity constraint:

$$e_3(B) \leq e_3(A). \quad (\text{C5})$$

In words, because agents underestimate their present bias, they think they will leave less effort for the last period than they end up leaving. We ignore this monotonicity constraint for now and verify that it holds later.

For each $\epsilon > 0$ small, consider a perturbation to the baseline efforts $\tilde{e}_2(B)$ and $\tilde{e}_3(B)$ that shifts effort from period 2 to period 3 according to self 2's preferences:

$$C(\tilde{e}_2(B)) = C(e_2(B)) + \epsilon, \quad C(\tilde{e}_3(B)) = C(e_3(B)) - \frac{\epsilon}{\beta\delta}.$$

By construction, this perturbation preserves (IC). Moreover, since the objective function evaluates costs from the perspective of self 1, this perturbation improves the objective. Thus, to minimize costs, the solution leaves as little effort as possible to the last period:

$$e_3(B) = 0. \quad (\text{C6})$$

It follows directly from (C6) that the monotonicity condition (C5) holds. Substituting back in (IC), we obtain

$$C(e_2(B)) = C(e_2(A)) + \beta\delta C(e_3(A)). \quad (\text{C7})$$

Substituting (C6) and (C7) in the objective function, we obtain

$$C(e_1) + \beta\delta C(e_2(A)) + (\beta\delta)^2 C(e_3(A)),$$

which is the cost of a time-consistent agent with discount factor $\beta\delta$.

Note how the argument above differs from the one in Lemma 2. The way to exploit naiveté in the consumption model is to postpone consumption in the baseline from period

2 to period 3. So, when deciding whether to consume now or to leave resources for the future, self 1 decides according to his long-run discount rate $\beta\delta^2$. Then, self 2 deviates from B to A, effectively bringing some consumption from $c_3(B)$ to period 2 (and reducing the consumption left to self 3). Since self 2 discounts period-3 consumption by $\beta\delta$, the rate between $u(c_1)$ and $u(c_2(A))$ is $\frac{\beta\delta^2}{\beta\delta} = \delta$, as shown in the auxiliary program in Lemma 2.

On the other hand, the way to exploit naiveté in the effort model is to require all effort in period 2, leaving zero effort for the future. Thus, self 1 decides how much effort to leave to the future according to his 1-period discount $\beta\delta$. Then, self 2 deviates from the baseline, leaving some effort for period 3. He decides how much to leave for period 3 also based on his 1-period discount $\beta\delta$. Therefore, the rate between $u(c_1)$ and $u(c_3(A))$ is $(\beta\delta)^2$.

APPENDIX D: SOPHISTICATED AGENTS

In Section 2, we characterized the equilibrium with either time-consistent or (partially) naive present-biased consumers. We now consider the case of sophisticated consumers, who correctly predict their future preferences ($\hat{\beta} = \beta$). We are interested in the asymptotic welfare of sophisticated consumers as the contracting horizon grows.

Recall that the welfare function does not discount future periods by the additional term β . Therefore, if consuming all resources in the first period maximizes welfare, the sophisticated agent must also consume all the resources in the first period. In this case, the equilibrium of the sophisticated agent trivially maximizes welfare. To rule out this uninteresting case, we proceed as in Section A.3 and assume that $\lim_{c \searrow 0} u'(c) = +\infty$. We will show that, when this is the case, the welfare loss from present bias of sophisticated agents does not vanish.

A sophisticated agent evaluates future consumption according to (1) with $\hat{\beta} = \beta$. Therefore, he fully understands that his future selves will behave like someone with the same time-consistency parameter as his current self. As with time-consistent consumers, since a sophisticated consumer agrees with the firm about his future preferences, there is no need to allow for options in the contract. Therefore, there is no loss of generality in restricting contracts to be vectors of state-dependent consumption. Because parties can commit to long-term contracts, any contract that is accepted by a firm must maximize the utility of the period-1 self subject to the zero-profits constraint. The equilibrium contract solves the following program:

$$\max_{\{c(s_t)\}} u(c(s_1)) + \beta E \left[\sum_{t=2}^T \delta^{t-1} u(c(s_t)) \right], \quad (\text{D1})$$

subject to the zero-profits constraint,

$$\sum_{t=1}^T E \left[\frac{w(s_t) - c(s_t)}{R^{t-1}} \right] = 0. \quad (\text{D2})$$

Let W_T^S denote the equilibrium welfare of the sophisticated consumer, which evaluates the consumption path according to the agent's long-run preferences (2), and recall that W_T^C is the welfare in the benchmark case of a time-consistent consumer. Since the time-consistent consumer maximizes welfare, the welfare loss from dynamic inconsistency cannot be negative:

$$W_T^C - W_T^S \geq 0.$$

We now show that unlike with partially naive agents, the consumption path of a sophisticated does not converge to the welfare-maximizing path as the contracting horizon grows. Therefore, the previous inequality is strict:

PROPOSITION 9: *Suppose u is bounded, $\delta < 1$, and $\lim_{c \searrow 0} u'(c) = +\infty$. Then, the welfare loss of a sophisticated consumer is uniformly bounded away from zero:*

$$\liminf_{T \rightarrow +\infty} (W_T^C - W_T^S) > 0.$$

Note that, in our model, the individual can consume in all periods, including when contracts are signed. If, instead, contracting occurred before consumption (say, at period 0), sophisticated consumers would commit to the ex ante optimal contract—see DellaVigna and Malmendier (2004), Heidhues and Köszegi (2010). The inefficiency with naive consumers, as well as the asymptotic efficiency result, remains unchanged if we add a contracting period with no consumption.

APPENDIX E: EQUILIBRIUM DEFINITION AND MIXED STRATEGIES

In this appendix, we present a formal definition of perception-perfect equilibria and show that the results in the paper generalize to mixed strategy equilibria.

For each state-dependent consumption $\{c(s_t)\}_{t \geq \tau}$, let

$$u(c(s_\tau)) + \beta E \left[\sum_{t > \tau} \delta^{t-\tau} u(c(s_t)) \middle| s_\tau \right] \quad (\text{E1})$$

denote self τ 's “experienced utility,” and let

$$u(c(s_\tau)) + \hat{\beta} E \left[\sum_{t > \tau} \delta^{t-\tau} u(c(s_t)) \middle| s_\tau \right] \quad (\text{E2})$$

denote self τ 's “perceived utility.”

As described in the text, it is without loss of generality to focus on contracts that offer a baseline (B) and an alternative (A) option in each period $t = 2, \dots, T - 1$. This is no longer true with mixed strategies. We generalize the definitions to allow for arbitrary message spaces when we consider mixed strategy equilibria.

Recall that we adopted the convention that each state describes all previous realization of uncertainty. A time- t history describes all actions by the consumer and all uncertainty realized until period t : $h_t = (\mathbf{c}, h^t, s_t)$, where \mathbf{c} is a contract offered by the time-1 self, $h^t \in \{A, B\}^{t-1}$ is an option history (which lists the options taken by all previous selves as defined in Section 2), and s_t is a state of the world at time t . Let \mathcal{H}_t denote the set of all possible time- t histories.

A pure strategy for the time-1 self is a consumption vector \mathbf{c} . A pure strategy for the firm is a mapping d from the space of possible consumption vectors to $\{0, 1\}$ specifying whether the firm accepts ($d = 1$) or rejects ($d = 0$) each consumption vector offered by the time-1 self. A pure strategy for self $t \in \{2, \dots, T - 1\}$ is a mapping from the time- t history to an option, that is, $\sigma_t: \mathcal{H}_t \rightarrow \{A, B\}$.³²

³²It is without loss of generality to focus on consumption vectors that do not offer any options to the time- T self, since it would be a dominant strategy for him to pick the one with the highest consumption (according to both the experienced and the perceived utility).

Before stating the equilibrium definition, we need to specify each player's payoffs. We start with the firm, which has correct beliefs. Let $\Pi(\mathbf{c}, \sigma_2, \dots, \sigma_{T-1}, \hat{d})$ denote the firm's expected profits from accepting ($\hat{d} = 1$) or rejecting ($\hat{d} = 0$) the consumption vector \mathbf{c} when selves $t > 1$ of the consumer play σ_t .

Since the consumer has incorrect beliefs about his future preferences, we need to distinguish between the actions that the consumer thinks he will choose and the actions that he ends up choosing. The agent's *perceived utility* determines what he thinks he will choose in the future, whereas the agent's *experienced utility* determines what he will end up choosing (see equations (E1) and (E2) in Appendix E):

- Let $U_1(\mathbf{c}, \sigma_2, \dots, \sigma_{T-1}, d)$ denote the time-1 self's expected experienced utility from offering contract \mathbf{c} if each future self $r > 1$ plays strategy σ_r and the firm plays d .
- For $t > 1$, let $U_t(\sigma_t, \dots, \sigma_{T-1} | h_t)$ denote the expected experienced utility of the time- t self conditional on history h_t when each self $r > t$ plays strategy σ_r .
- Let $\hat{U}_t(\sigma_t, \dots, \sigma_{T-1} | h_t)$ denote the expected perceived utility of the time- t self conditional on history h_t when each self $r > t$ plays strategy σ_r .

DEFINITION 1: A perception-perfect equilibrium is a consumption vector \mathbf{c} , a pair of strategies $(\sigma_2, \dots, \sigma_{T-1})$ and $(\hat{\sigma}_2, \dots, \hat{\sigma}_{T-1})$, and an acceptance decision d such that:

- \mathbf{c} maximizes self 1's expected experienced utility under the assumption that his future selves use strategy $\hat{\sigma}_r$ and the firm uses strategy d :

$$U_1(\mathbf{c}, \hat{\sigma}_2, \dots, \hat{\sigma}_{T-1}, d) \geq U_1(\mathbf{c}', \hat{\sigma}_2, \dots, \hat{\sigma}_{T-1}, d), \quad \forall \mathbf{c}'.$$

- For all $t > 1$ and all h_t , $\sigma_t(h_t)$ maximizes self- t 's expected experienced utility under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$:

$$U_t(\sigma_t, \hat{\sigma}_{t+1}, \dots, \hat{\sigma}_{T-1} | h_t) \geq U_t(\sigma'_t, \hat{\sigma}_{t+1}, \dots, \hat{\sigma}_{T-1} | h_t), \quad \forall \sigma'_t.$$

- For all $t > 1$ and all h_t , $\hat{\sigma}_t(h_t)$ maximizes the consumer's time- t expected perceived utility under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$:

$$\hat{U}_t(\hat{\sigma}_t, \hat{\sigma}_{t+1}, \dots, \hat{\sigma}_{T-1} | h_t) \geq \hat{U}_t(\sigma'_t, \hat{\sigma}_{t+1}, \dots, \hat{\sigma}_{T-1} | h_t), \quad \forall \sigma'_t.$$

- For all \mathbf{c} , $d(\mathbf{c})$ maximizes the firm's expected discounted profits under the assumption that the consumer uses strategies σ_t for all t :

$$\Pi(\mathbf{c}, \sigma_2, \dots, \sigma_{T-1}, d) \geq \Pi(\mathbf{c}, \sigma_2, \dots, \sigma_{T-1}, d'), \quad \forall d'.$$

We now generalize the definition of perception-perfect equilibrium to allow for mixed strategies. In this case, we need to work with more general message spaces, since the restriction to two possible messages is no longer without loss of generality.

Let M_2 be a non-empty, compact space of possible messages in period 2 with generic element m_2 . For each $t \in \{3, \dots, T-1\}$, let $M_t(m_2, \dots, m_{t-1})$ be a non-empty, compact space of possible messages in period t conditional on previous messages (m_2, \dots, m_{t-1}) . Let \mathbf{M} denote the space of all possible messages in all periods.³³ A consumption vector (or "contract") \mathbf{c} specifies, for each period, the consumption conditional on all messages up to t and all realized uncertainty: $c(m_2, \dots, m_{t-1}, m_t, s_t)$. A period- t history

³³That is, $\mathbf{M} \equiv \{M_1, \dots, M_{T-1}(m_1, \dots, m_{T-2}) : m_1 \in M_1, \dots, (m_1, \dots, m_{T-1}) \in M_1 \times \dots \times M_{T-1}(m_1, \dots, m_{T-2})\}$.

$h_t = (\mathbf{c}, \mathbf{M}, m_2, \dots, m_{t-1}, s_t)$ consists of a consumption vector and a message space offered at time 1, the messages sent in all previous periods, and the state of the world at t describing all realized uncertainty.

With a slight abuse of notation, we now allow σ to be a mixed strategy as well. A mixed strategy for the time-1 self σ_1 is a distribution over (compact and non-empty) message spaces and contracts. A mixed strategy for the firm σ_{firm} is a distribution over acceptance decisions for each contract and message space offered by the time-1 self. A mixed strategy for self $t \in \{2, \dots, T-1\}$ specifies, for each period- t history, a distribution over messages: $\sigma_t(\mathbf{c}, m_2, \dots, m_{t-1}, s_t) \in \Delta(M_t(m_2, \dots, m_{t-1}))$.

We now extend the payoffs to allow for mixed strategies:

- Let $\Pi(\sigma_1, \dots, \sigma_{T-1}, \sigma_{\text{firm}})$ denote the firm's expected profits from playing σ_{firm} when each consumer self plays strategy σ_1 .
- Let $U_1(\sigma_1, \dots, \sigma_{T-1}, \sigma_{\text{firm}})$ denote the expected experienced utility (E1) of the time-1 self from playing strategy σ_1 if each future self $r > 1$ plays σ_r and the firm plays σ_{firm} .
- For $t > 1$, let $U_t(\sigma_t, \dots, \sigma_{T-1} | h_t)$ denote the expected experienced utility (E1) of the time- t self conditional on history h_t when each self $r > t$ plays strategy σ_r .
- For $t > 1$, let $\hat{U}_t(\sigma_t, \dots, \sigma_{T-1} | h_t)$ denote the expected perceived utility (E2) of the time- t self conditional on history h_t when each self $r > t$ plays strategy σ_r .

We can now state the equilibrium definition:

DEFINITION 2: A perception-perfect equilibrium in mixed strategies is a pair of strategies for the consumer $(\mathbf{c}, \sigma_2, \dots, \sigma_{T-1})$ and $(\hat{\sigma}_2, \dots, \hat{\sigma}_{T-1})$ and a strategy for the firm σ_{firm} such that:

- σ_1 maximizes self 1's expected experienced utility under the assumption that his future selves use strategy $\hat{\sigma}_r$ and the firm uses strategy σ_{firm} :

$$U_1(\sigma_1, \hat{\sigma}_2, \dots, \hat{\sigma}_{T-1}, \sigma_{\text{firm}}) \geq U_1((\mathbf{c}', \mathbf{M}'), \hat{\sigma}_2, \dots, \hat{\sigma}_{T-1}, \sigma_{\text{firm}}), \quad \forall \mathbf{c}', \mathbf{M}'.$$

- For all $t > 1$ and all h_t , $\sigma_t(h_t)$ maximizes self- t 's expected experienced utility under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$:

$$U_t(\sigma_t, \hat{\sigma}_{t+1}, \dots, \hat{\sigma}_{T-1} | h_t) \geq U_t(m_t, \hat{\sigma}_{t+1}, \dots, \hat{\sigma}_{T-1} | h_t), \quad \forall m_t.$$

- For all $t > 1$ and all h_t , $\hat{\sigma}_t(h_t)$ maximizes the consumer's time- t expected perceived utility under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$:

$$\hat{U}_t(\hat{\sigma}_t, \hat{\sigma}_{t+1}, \dots, \hat{\sigma}_{T-1} | h_t) \geq \hat{U}_t(m_t, \hat{\sigma}_{t+1}, \dots, \hat{\sigma}_{T-1} | h_t), \quad \forall m_t.$$

- For all \mathbf{c} , $\sigma_{\text{firm}}(\mathbf{c})$ maximizes the firm's expected discounted profits under the assumption that the consumer uses strategies σ_t in periods $t > 1$:

$$\Pi(\mathbf{c}, \mathbf{M}, \sigma_2, \dots, \sigma_{T-1}, \sigma_{\text{firm}}) \geq \Pi(\mathbf{c}, \mathbf{M}, \sigma_2, \dots, \sigma_{T-1}, d'), \quad \forall d' = 0, 1.$$

We can now establish that our restriction to pure strategies in the text was without loss of generality. The *equilibrium program* is

$$\max_{\{c(s_t, h^t)\}} u(c(s_1)) + \beta E \left[\sum_{t=2}^T \delta^{t-1} u(c(s_t, \hat{\sigma}_2, \hat{\sigma}_3, \dots, \hat{\sigma}_{T-1})) \right], \quad (\text{E3})$$

subject to

$$\sum_{t=1}^T E \left[\frac{w(s_t) - c(s_t, \sigma_2, \sigma_3, \dots, \sigma_{T-1})}{R^{t-1}} \right] = 0, \quad (\text{Zero Profits})$$

$$\begin{aligned} & u(c(s_\tau, (h^{\tau-1}, \hat{m}_\tau))) + \hat{\beta} E \left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, \hat{m}_\tau, \hat{\sigma}_{\tau+1}, \dots, \hat{\sigma}_{T-1}))) \middle| s_\tau \right] \\ & \geq u(c(s_\tau, (h^{\tau-1}, m'_\tau))) + \hat{\beta} E \left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, m'_\tau, \hat{\sigma}_{\tau+1}, \dots, \hat{\sigma}_{T-1}))) \middle| s_\tau \right], \\ & \forall \hat{m}_\tau \in \text{supp}(\hat{\sigma}_\tau), m'_\tau \in M_\tau \end{aligned} \quad (\text{PC})$$

and

$$\begin{aligned} & u(c(s_\tau, (h^{\tau-1}, m_\tau))) + \beta E \left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, m_\tau, \hat{\sigma}_{\tau+1}, \dots, \hat{\sigma}_{T-1}))) \middle| s_\tau \right] \\ & \geq u(c(s_\tau, (h^{\tau-1}, m'_\tau))) + \beta E \left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, m'_\tau, \hat{\sigma}_{\tau+1}, \dots, \hat{\sigma}_{T-1}))) \middle| s_\tau \right], \\ & \forall m_\tau \in \text{supp}(\sigma_\tau), m'_\tau \in M_\tau. \end{aligned} \quad (\text{IC})$$

The next lemma establishes that the consumption path for time-inconsistent agents still coincides with the solution of the auxiliary program when we allow for mixed strategies:

LEMMA 9: *In any perception-perfect equilibrium in mixed strategies, the consumption path solves the auxiliary program (7).*

The proof is in the supplementary appendix. To illustrate it, consider the case in which $T = 3$ and the consumer has a constant income w in each period. Suppose self 1 believes that self 2 will pick option B_1 with probability θ_1 and option B_2 with probability θ_2 . (PC) states that

$$u(c_2(B_1)) + \hat{\beta} \delta u(c_3(B_1)) = u(c_2(B_2)) + \hat{\beta} \delta u(c_3(B_2)) \geq u(c_2(A)) + \hat{\beta} \delta u(c_3(A)),$$

whereas (IC) requires

$$\begin{aligned} u(c_2(A)) + \beta \delta u(c_3(A)) & \geq u(c_2(B_1)) + \beta \delta u(c_3(B_1)), \\ u(c_2(A)) + \beta \delta u(c_3(A)) & \geq u(c_2(B_2)) + \beta \delta u(c_3(B_2)). \end{aligned}$$

So self 1's perceived utility is

$$u(c_1) + \beta \delta (\theta_1 u(c_2(B_1)) + \theta_2 u(c_2(B_2))) + \beta \delta^2 (\theta_1 u(c_3(B_1)) + \theta_2 u(c_3(B_2))).$$

We claim that both ICs must bind. First, note that at least one of them must bind (otherwise, we can raise $c_3(B_1)$ and $c_3(B_2)$ without affecting any other constraints). Suppose the IC associated with B_1 binds but not the one associated with B_2 :

$$u(c_2(B_1)) + \beta \delta u(c_3(B_1)) = u(c_2(A)) + \beta \delta u(c_3(A)) > u(c_2(B_2)) + \beta \delta u(c_3(B_2)).$$

Recall that self 1 perceives that self 2 will mix between option B_1 and option B_2 , so self 1 believes that self 2 must be indifferent:

$$u(c_2(B_1)) + \hat{\beta}\delta u(c_3(B_1)) = u(c_2(B_2)) + \hat{\beta}\delta u(c_3(B_2)).$$

It follows that $c_2(B_1) > c_2(B_2)$, $c_3(B_1) < c_3(B_2)$. Together with the perceived-choice constraints, these inequalities imply that

$$u(c_2(B_1)) + \delta u(c_3(B_1)) < u(c_2(B_2)) + \delta u(c_3(B_2)).$$

Consider an alternative contract that sets the consumption associated with option B_1 equal to the consumption associated with option B_2 . This contract strictly increases self 1's perceived utility, a contradiction to the optimality of the original contract. So both IC constraints are binding, implying that $c_2(B_1) = c_2(B_2)$ and $c_3(B_1) = c_3(B_2)$. Similarly to the proof of Lemma 2, $c_2(B_1) = c_2(B_2) = 0$. Substituting them back to the objective function leads to the auxiliary program

$$u(c_1) + \delta u(c_2(A)) + \beta\delta^2 u(c_3(A)).$$

On the other hand, suppose the firm believes that the consumer will randomize between two alternative options, A_1 and A_2 , with probabilities θ_1 and θ_2 , respectively. Following the same steps in the proof of Lemma 2, we obtain the auxiliary program:

$$u(c_1) + \delta[\theta_1 u(c_2(A_1)) + \theta_2 u(c_2(A_2))] + \beta\delta^2[\theta_1 u(c_3(A_1)) + \theta_2 u(c_3(A_2))],$$

subject to the zero-profit condition

$$c_1 + \frac{\theta_1 c_2(A_1) + \theta_2 c_2(A_2)}{R} + \frac{\theta_1 c_3(A_1) + \theta_2 c_3(A_2)}{R^2} = w + \frac{w}{R} + \frac{w}{R^2}.$$

If options A_1 and A_2 are different, by Jensen's inequality and the strict concavity of $u(\cdot)$, merging these two options A_1 and A_2 would strictly increase self 1's payoff.

REFERENCES

- BANKS, J. S., AND J. SOBEL (1987): "Equilibrium Selection in Signaling Games," *Econometrica*, 55 (3), 647–661. [3]
 DELLAVIGNA, S., AND U. MALMENDIER (2004): "Contract Design and Self-Control: Theory and Evidence," *Quarterly Journal of Economics*, 119 (2), 353–402. [9]
 HEIDHUES, P., AND B. KŐSZEGI (2010): "Exploiting Naivete About Self-Control in the Credit Market," *American Economic Review*, 100 (5), 2279–2303. [9]

Co-editor Alessandro Lizzeri handled this manuscript.

Manuscript received 6 March, 2019; final version accepted 23 September, 2020; available online 25 September, 2020.