

Addiction and Cue-Triggered Decision Processes:
Complete Appendix

B. Douglas Bernheim
Department of Economics
Stanford University
and
NBER

Antonio Rangel
Department of Economics
Stanford University
and
NBER

July 2004

For convenience, we replicate here the main value equation from the text:

$$V_s(\theta) = \max_{(a,x) \in \{(E,1), (E,0), (A,0), (R,0)\}} u_s^a + \sigma_s^{a,x} b_s^a + \delta [(1 - \sigma_s^{a,x}) V_{\max\{1, s-1\}}(\theta) + \sigma_s^{a,x} V_{\min\{S, s+1\}}(\theta)] \quad (1)$$

Proof of Proposition 1

Proof of parts (i-a) and (i-b). We begin the proof with two lemmas.

Lemma 1: Consider θ' and θ'' such that: (1) $\theta'_k \neq \theta''_k$, (2) $\theta'_i = \theta''_i$ for $i \neq k$, and (3) $V_s(\theta') \geq V_s(\theta'')$ for all s . Then:

$$(a) \text{ for all } j < k, V_j(\theta') - V_j(\theta'') \leq V_{j+1}(\theta') - V_{j+1}(\theta''),$$

$$(b) \text{ for all } j > k, V_j(\theta') - V_j(\theta'') \leq V_{j-1}(\theta') - V_{j-1}(\theta'').$$

Proof: We provide a proof of (a). The argument for (b) is symmetric. The proof uses the following notation:

$$\lambda_s^{a,1}(\theta) = u_s^a + b_s^a + \delta V_{\min\{S, s+1\}}(\theta) \text{ for } a = E, A \quad (2)$$

$$\lambda_s^{a,0}(\theta) = u_s^a + p_s^a b_s^a + (1 - p_s^a) \delta V_{\max\{1, s-1\}}(\theta) + p_s^a \delta V_{\min\{S, s+1\}}(\theta) \text{ for } a = E, A \quad (3)$$

$$\lambda_s^{R,0}(\theta) = u_s^R - c_s + \delta V_{\max\{1, s-1\}}(\theta) \quad (4)$$

We claim that for all $j < k$, $V_j(\theta') - V_j(\theta'') \leq V_{j+1}(\theta') - V_{j+1}(\theta'')$. Consider first the case $j = 1$. (The case $j = 0$ is almost identical and thus is omitted).¹

$$\begin{aligned} V_1(\theta') - V_1(\theta'') &= \max_{(a,x) \in \{(E,1), (E,0), (A,0), (R,0)\}} \lambda_1^{a,x}(\theta') - \max_{(a,x) \in \{(E,1), (E,0), (A,0), (R,0)\}} \lambda_1^{a,x}(\theta'') \\ &\leq \max_{(a,x) \in \{(E,1), (E,0), (A,0), (R,0)\}} (\lambda_1^{a,x}(\theta') - \lambda_1^{a,x}(\theta'')) \\ &= \delta \max\{V_2(\theta') - V_2(\theta''), (1 - p_1^E)(V_1(\theta') - V_1(\theta'')) + p_1^E(V_2(\theta') - V_2(\theta'')), \\ &\quad (1 - p_1^A)(V_1(\theta') - V_1(\theta'')) + p_1^A(V_2(\theta') - V_2(\theta'')), V_1(\theta') - V_1(\theta'')\} \end{aligned}$$

(where p_1^a corresponds to both θ' and θ''). Consider the last expression. Given the second and third terms are each weighted averages of the first and fourth terms, either the first or fourth term must be the maximand. If the fourth term is the maximand then, since $\delta \in (0, 1)$, $V_1(\theta') - V_1(\theta'') = 0 \leq V_2(\theta') - V_2(\theta'')$ (the last inequality follows from condition (3) of the lemma). If the first term is the maximand, the claim trivially holds.

¹The argument makes use of the fact that, for any eight real numbers z_1, \dots, z_8 , $\max_{i \in \{1, 2, 3, 4\}} z_i - \max_{i \in \{5, 6, 7, 8\}} z_i \leq \max_{i \in \{1, 2, 3, 4\}} \{z_i - z_{i+4}\}$.

Now consider the following induction step. We show that for all $j < k$,

$$V_{j-1}(\theta') - V_{j-1}(\theta'') \leq V_j(\theta') - V_j(\theta'') \Rightarrow V_j(\theta') - V_j(\theta'') \leq V_{j+1}(\theta') - V_{j+1}(\theta'').$$

Arguing as above, we have

$$\begin{aligned} V_j(\theta') - V_j(\theta'') &\leq \delta \max\{V_{j+1}(\theta') - V_{j+1}(\theta''), \\ &\quad (1 - p_j^E)(V_{j-1}(\theta') - V_{j-1}(\theta'')) + p_j^E(V_{j+1}(\theta') - V_{j+1}(\theta'')), \\ &\quad (1 - p_j^A)(V_{j-1}(\theta') - V_{j-1}(\theta'')) + p_j^A(V_{j+1}(\theta') - V_{j+1}(\theta'')), V_{j-1}(\theta') - V_{j-1}(\theta'')\} \end{aligned} \quad (5)$$

(where p_j^a corresponds to both θ' and θ''). As before, there are two possible cases. If the fourth term is the maximand we obtain

$$V_j(\theta') - V_j(\theta'') \leq \delta(V_{j-1}(\theta') - V_{j-1}(\theta'')) \leq \delta(V_j(\theta') - V_j(\theta''));$$

where the last inequality follows from the induction hypothesis. This implies that $V_j(\theta') - V_j(\theta'') = 0 \leq V_{j+1}(\theta') - V_{j+1}(\theta'')$. If the first term is the maximand, the claim trivially holds. (Note that this establishes the claim only for $j < k$; (5) does not hold for $j = k$ since $\theta'_k \neq \theta''_k$). Q.E.D.

Lemma 2: Consider θ' and θ'' such that: (1) $\theta'_k \neq \theta''_k$, (2) $\theta'_i = \theta''_i$ for $i \neq k$, and (3) $V_s(\theta') \geq V_s(\theta'')$ for all s . Then:

- (a) For $j < k$, the disposition to use in state j is weakly higher with θ' than with θ'' ,
- (b) For $j > k$, the disposition to use in state j is weakly lower with θ' than with θ'' .

Proof: Consider any θ' and θ'' that differ only with respect to state k , and assume that valuation in all states is weakly higher with θ' than with θ'' . Lemma 1 tells us that, for $j < k$,

$$V_{\max\{1, j-1\}}(\theta') - V_{\max\{1, j-1\}}(\theta'') \leq V_{j+1}(\theta') - V_{j+1}(\theta'').$$

(The case $j = 0$ is identical and thus is omitted). Rearranging this expression yields:

$$\Delta V_j(\theta') \leq \Delta V_j(\theta'') \quad (6)$$

where

$$\Delta V_j(\theta) \equiv V_{\max\{1, j-1\}}(\theta) - V_{\min\{S, j+1\}}(\theta).$$

Define

$$\mu_j^{E1}(\theta) \equiv \frac{b_j^E}{\delta},$$

$$\begin{aligned}
\mu_j^{A0}(\theta) &\equiv \frac{(u_j^E - u_j^A) + (p_j^E b_j^E - p_j^A b_j^A)}{\delta(p_j^E - p_j^A)}, \\
\mu_j^{R0}(\theta) &\equiv \frac{b_j^E}{\delta} + \frac{u_j^E - u_j^R}{\delta p_j^E}, \\
\mu_j^{AR0}(\theta) &\equiv \frac{b_j^A}{\delta} + \frac{u_j^A - u_j^R}{\delta p_j^A}.
\end{aligned}$$

Simple algebraic manipulation of (1) reveals that

$$(E, 1) \text{ is optimal iff } \Delta V_j(\theta) \leq \mu_j^{E1}(\theta) \quad (7)$$

$$(E, 0) \text{ is optimal iff } \Delta V_j(\theta) \in (\mu_j^{E1}(\theta), \min\{\mu_j^{A0}(\theta), \mu_j^{R0}(\theta)\}) \quad (8)$$

$$(A, 0) \text{ is optimal iff } \mu_j^{A0}(\theta) \leq \mu_j^{R0}(\theta) \text{ and } \Delta V_j(\theta) \in (\mu_j^{A0}(\theta), \mu_j^{AR0}(\theta)) \quad (9)$$

$$\begin{aligned}
(R, 0) \text{ is optimal iff either } \mu_j^{A0}(\theta) &\leq \mu_j^{R0}(\theta) \text{ and } \Delta V_j(\theta) \geq \mu_j^{AR0}(\theta), \\
\text{or } \mu_j^{A0}(\theta) &\geq \mu_j^{R0}(\theta) \text{ and } \Delta V_j(\theta) \geq \mu_j^{R0}(\theta)
\end{aligned} \quad (10)$$

Note that $\mu_j^z(\theta') = \mu_j^z(\theta'')$ for $z = E1, E0, A0$, and $R0$. From (6) and (7) through (10), it then follows that the disposition to use in state j is weakly higher with θ' than with θ'' . A symmetric argument for $j > k$ completes the proof. Q.E.D.

It is easy to verify that $V_s(\theta)$ is weakly increasing in u_k^a and b_k^a , and weakly decreasing in p_k^a . Combining this with Lemma 2 completes the proof of parts (i-a) and (i-b).

Proof of part (i-c). Consider two parameter vectors, $\underline{\theta}$ and $\bar{\theta}$, such that $\bar{b}_j^E > \underline{b}_j^E$ with all other components equal, or $\bar{p}_j^E < \underline{p}_j^E$ with all other components equal. We argue, in three steps, that the disposition to use is weakly higher with $\bar{\theta}$ than with $\underline{\theta}$.

Step 1: (a) If $(E, 1)$ is optimal in state j with $\underline{\theta}$, then it is optimal in state j with $\bar{\theta}$, and (b) if $(E, 1)$ is the unique optimal choice in state j with $\underline{\theta}$, then it is the unique optimal choice in state j with $\bar{\theta}$.

For $\underline{\theta}$ and $\bar{\theta}$ with $\bar{p}_j^E < \underline{p}_j^E$ and all other components equal, (a) and (b) follow from part (ii) of the proposition, proven below. Here, we consider $\underline{\theta}$ and $\bar{\theta}$ with $\bar{b}_j^E > \underline{b}_j^E$ and all other components equal.

Consider any optimal decision function $\bar{\chi}_s : \{0, 1, \dots, S\} \rightarrow \{(E, 1), (E, 0), (A, 0), (R, 0)\}$ for the parameter vector $\bar{\theta}$. Imagine that the DM follows the optimal decision rule $\bar{\chi}_s$, and that he starts from state $j - 1$ at age 0. Let g_t indicate the probability of reaching state $j + 1$ for the first time at age t (note that $g_1 = 0$, and that $g_s = 0$ for all s when $\bar{\chi}_i = (R, 0)$ for $i = j, j - 1$). Let G_t indicate the expected discounted payoff for ages 0 through $t - 2$, conditional upon reaching state $j + 1$ for the first time at age t . If there is a positive probability that state $j + 1$ will never be reached, let G_∞ denote the expected payoff for all periods conditional on this event (otherwise let $G_\infty = 0$). Note

that G_t and G_∞ are the same regardless of whether one evaluates payoffs under $\bar{\theta}$ and $\underline{\theta}$. Note that we can write:

$$V_{j-1}(\bar{\theta}) = \sum_{t=2}^{\infty} \left[\frac{G_t}{\delta^{t-1}} + \bar{u}_j^E + \bar{b}_j^E + \delta V_{j+1}(\bar{\theta}) \right] g_t \delta^{t-1} + \left[1 - \sum_{t=2}^{\infty} g_t \right] G_\infty$$

Since the DM has the option to follow $\bar{\chi}_s$ for $s < j+1$, we know that

$$V_{j-1}(\underline{\theta}) \geq \sum_{t=2}^{\infty} \left[\frac{G_t}{\delta^{t-1}} + \underline{u}_j^E + \underline{b}_j^E + \delta V_{j+1}(\underline{\theta}) \right] g_t \delta^{t-1} + \left[1 - \sum_{t=2}^{\infty} g_t \right] G_\infty$$

Since $\bar{u}_j^E = \underline{u}_j^E$, we have

$$V_{j-1}(\bar{\theta}) - V_{j-1}(\underline{\theta}) \leq \sum_{t=2}^{\infty} \left[\bar{b}_j^E - \underline{b}_j^E + \delta(V_{j+1}(\bar{\theta}) - V_{j+1}(\underline{\theta})) \right] g_t \delta^{t-1}$$

Since $\sum_{t=2}^{\infty} g_t \delta^{t-1} < 1$ (and since $\bar{b}_j^E > \underline{b}_j^E$ implies $V_{j+1}(\bar{\theta}) \geq V_{j+1}(\underline{\theta})$), it follows that

$$V_{j-1}(\bar{\theta}) - V_{j-1}(\underline{\theta}) < \bar{b}_j^E - \underline{b}_j^E + \delta(V_{j+1}(\bar{\theta}) - V_{j+1}(\underline{\theta}))$$

Consequently,

$$\begin{aligned} \Delta V_j(\bar{\theta}) - \Delta V_j(\underline{\theta}) &= [V_{j-1}(\bar{\theta}) - V_{j+1}(\bar{\theta})] - [V_{j-1}(\underline{\theta}) - V_{j+1}(\underline{\theta})] \\ &= [V_{j-1}(\bar{\theta}) - V_{j-1}(\underline{\theta})] - [V_{j+1}(\bar{\theta}) - V_{j+1}(\underline{\theta})] \\ &< (\bar{b}_j^E - \underline{b}_j^E) - (1 - \delta)(V_{j+1}(\bar{\theta}) - V_{j+1}(\underline{\theta})) \\ &< \frac{\bar{b}_j^E - \underline{b}_j^E}{\delta} \end{aligned}$$

(where, in the last step, we have used $\delta < 1$ along with the fact that $V_{j+1}(\bar{\theta}) \geq V_{j+1}(\underline{\theta})$). Since, by assumption, $(E, 1)$ is optimal in state j with $\underline{\theta}$, we know that $\Delta V_j(\underline{\theta}) \leq \frac{\bar{b}_j^E}{\delta}$. Consequently, we have

$$\begin{aligned} \Delta V_j(\bar{\theta}) &< \Delta V_j(\underline{\theta}) + \frac{\bar{b}_j^E - \underline{b}_j^E}{\delta} \\ &\leq \frac{\bar{b}_j^E}{\delta} + \frac{\bar{b}_j^E - \underline{b}_j^E}{\delta} = \frac{\bar{b}_j^E}{\delta} \end{aligned}$$

By (7) this implies $(E, 1)$ is the unique optimal choice in state j with $\bar{\theta}$.

Step 2: (a) If neither $(E, 1)$ nor $(E, 0)$ are optimal choices in state j for $\bar{\theta}$, then the sets of optimal state j choices are identical with $\bar{\theta}$ and $\underline{\theta}$; (b) if either $(A, 0)$ or $(R, 0)$ is optimal in state j with $\bar{\theta}$, it is also optimal with $\underline{\theta}$.

Given that $V_s(\bar{\theta})$ satisfies (1) when $u_s^a = \bar{u}_s^a$ and $b_s^a = \bar{b}_s^a$, and $\sigma_s^{a,x} = \bar{\sigma}_s^{a,x}$ with either $(A, 0)$ or $(R, 0)$ as the maximizing choice in state j , $V_s(\underline{\theta}) = V_s(\bar{\theta})$ plainly satisfies (1) when $u_s^a = \underline{u}_s^a$, $b_s^a = \underline{b}_s^a$, and $\sigma_s^{a,x} = \bar{\sigma}_s^{a,x}$. Claims (a) and (b) follow directly.

Step 3. The disposition to use is weakly higher with $\bar{\theta}$ than with $\underline{\theta}$. First we show that the maximum disposition to use is weakly higher with $\bar{\theta}$ than with $\underline{\theta}$. If

the maximum disposition to use is $(E, 1)$ with $\bar{\theta}$, the claim is immediate; if it's $(E, 0)$ with $\bar{\theta}$, part (a) of step 1 tells us it can't be $(E, 1)$ with $\underline{\theta}$; if it's $(A, 0)$ or $(R, 0)$ with $\bar{\theta}$, part (a) of step 2 tells us it's the same with $\underline{\theta}$. Next we show that the minimum disposition to use is weakly higher with $\bar{\theta}$ than with $\underline{\theta}$. If the minimum disposition to use is $(E, 1)$ with $\bar{\theta}$, the claim is immediate; if it's $(E, 0)$ with $\bar{\theta}$, part (b) of step 1 tells us it can't be $(E, 1)$ with $\underline{\theta}$; if it's $(A, 0)$ or $(R, 0)$ with $\bar{\theta}$, part (b) of step 2 tells us it can't be greater with $\underline{\theta}$.

Now consider two parameter vectors, $\underline{\theta}$ and $\bar{\theta}$, such that $\bar{u}_j^R < \underline{u}_j^R$ with all other components equal. We claim that if something other than $(R, 0)$ is optimal in state j with $\underline{\theta}$, then it is also optimal in state j with $\bar{\theta}$ (from which it follows that the maximum disposition to use cannot be higher with $\underline{\theta}$); moreover, if $(R, 0)$ is not optimal in state j with $\underline{\theta}$, then the sets of optimal state j choices are identical with $\bar{\theta}$ and $\underline{\theta}$ (from which it follows that the minimum disposition to use cannot be higher with $\underline{\theta}$). Analogously to step 2, these conclusions follow from the fact that $V_s(\underline{\theta})$ continues to satisfy (1) when the parameter vector is $\bar{\theta}$.

Proof of part (ii). Suppose $\bar{\theta}$ coincides with $\underline{\theta}$ except for $p_j^E, p_j^A, u_j^A, u_j^R$, and/or b_j^A (subject to the restrictions imposed by Assumptions 1 and 2). We claim that, if $(E, 1)$ is optimal in state j for $\underline{\theta}$, it is also optimal in state j for $\bar{\theta}$; moreover, if $(E, 1)$ is the unique optimum in the first instance it is also the unique optimum in the second instance. Part (ii) follows directly from these claims. By construction, $V_s(\underline{\theta})$ satisfies (1) for $\theta = \underline{\theta}$. We argue that, if $(E, 1)$ is optimal in state j for $\underline{\theta}$, then $V_s(\bar{\theta}) = V_s(\underline{\theta})$ also satisfies (1) for $\theta = \bar{\theta}$. Under the hypothesis that $V_s(\bar{\theta}) = V_s(\underline{\theta})$, we have $\lambda_s^{a,x}(\bar{\theta}) = \lambda_s^{a,x}(\underline{\theta})$ for all $s \neq j$, and for $(s, a, x) = (j, E, 1)$. Thus, (1) is satisfied for all $s \neq j$. Since $(E, 1)$ is optimal for state j with $\underline{\theta}$, we know that $(E, 1)$ solves (1) for state j with $\theta = \underline{\theta}$, which is equivalent to $\Delta V_j(\underline{\theta}) \leq \frac{b_j^E}{\delta}$. But then, under our hypothesis, $\Delta V_j(\bar{\theta}) \leq \frac{b_j^E}{\delta}$ as well, so $(E, 1)$ remains a maximizer for state j with $\theta = \bar{\theta}$. Since $\lambda_s^{E,1}(\underline{\theta}) = \lambda_s^{E,1}(\bar{\theta})$, the maximized value of (1) is unchanged with $\theta = \bar{\theta}$. Accordingly, $V_s(\bar{\theta}) = V_s(\underline{\theta})$ is the maximized value function when $\theta = \bar{\theta}$, and $(E, 1)$ is an optimal choice in state j . If $(E, 1)$ is the unique maximizer with $\underline{\theta}$, then $\Delta V_j(\bar{\theta}) = \Delta V_j(\underline{\theta}) < \frac{b_j^E}{\delta}$, so it is also the unique maximizer with $\bar{\theta}$. Q.E.D.

Proof of Proposition 2

To avoid repetition, we organize the proof by aspects of behavior, rather than by groups of parameters (as in the statement of the proposition). We establish each element of the proof by comparing behavior for pairs of parameter vectors, labeled $\underline{\theta}$ and $\bar{\theta}$.

Choice in state 0. Suppose $\underline{\eta}_s \leq \bar{\eta}_s$ for $\eta = p^E, p^A, u^A, u^R, b^E$, or b^A (with all other parameters fixed). We can divide this change into two components: a weak

increase in η_0 , and a weak increase in η_s for $s > 0$. First consider the weak increase in η_s for $s > 0$. The impact on the disposition to use in state 0 follows from Proposition 1, part (i-a). Next consider the weak increase in η_0 . For $\eta = p^a$, $\underline{p}_0^a = \bar{p}_0^a = 0$, so there is no change to consider. For the case of $\eta = u^A$, u^R , or b^A , the increase is irrelevant (the DM only selects $(E, 1)$ or $(E, 0)$ in state 0). For the case of $\eta = b^E$, an increase in b_0^E increases the disposition to use in state 0, which reinforces the effect of increasing b_s^E for $s > 0$.

Choice in state 1. Suppose $\underline{p}_s^E \leq \bar{p}_s^E$ for all s (with all other parameters fixed). We can divide the difference between $\underline{\theta}$ and $\bar{\theta}$ into two components: a weak increase in p_1^E , and a weak increase in p_s^E for $s > 1$. The first change weakly reduces the disposition to use in state 1 (Proposition 1, part (i-c)), as does the second change (Proposition 1, part (i-a)). The argument for a general decrease in b_s^E is identical, but requires the added observation that b_0^E is irrelevant (since state 0 is unattainable from state 1).

Choice in state S . Suppose $\underline{u}_s^R \leq \bar{u}_s^R$ for all s (with all other parameters fixed). We can divide the difference between $\underline{\theta}$ and $\bar{\theta}$ into two components: an increase in u_S^R , and an increase in u_s^R for $s < S$. The first change reduces the disposition to use in state S (Proposition 1, part (i-c)), as does the second change (Proposition 1, part (i-b)).

Choice in intermediate states. Suppose $\underline{u}_s^R \leq \bar{u}_s^R$ for all s (with all other parameters fixed). With $\bar{\theta}$, let \bar{s}^* denote the first state in which $(R, 0)$ is an optimal choice. Consider moving from $\bar{\theta}$ to $\underline{\theta}$ in two steps. (1) One state at a time, change from \bar{u}_s^R to \underline{u}_s^R for $s < \bar{s}^*$. Since $(R, 0)$ is not chosen in any of these states with $\bar{\theta}$, this does not change the value function, and the same choices remain optimal in every state. (2) Change from \bar{u}_s^R to \underline{u}_s^R for $s \geq \bar{s}^*$. This weakly increases the disposition to use in state $s < \bar{s}^*$ (Proposition 1, part (i-a)). Thus, the disposition to use weakly increases in all states $s < \bar{s}^*$.

First intentional use interval. Suppose $\underline{p}_s^E \geq \bar{p}_s^E$ for all s (with all other parameters fixed). With $\underline{\theta}$, let \underline{s}^1 denote the first state (other than 0) in which $(E, 1)$ is not an optimal choice. Let $\underline{s}_0^1 \leq \underline{s}^1$ denote the first state (other than 0) in which something other than $(E, 1)$ is an optimal choice. Consider moving from $\underline{\theta}$ to $\bar{\theta}$ in two steps. (1) Change from \underline{p}_s^E to \bar{p}_s^E for $s < \underline{s}^1$. Since $(E, 1)$ is initially optimal for all such states, this leaves all optimal choices unchanged (Proposition 1, part (ii)). (2) Change from \underline{p}_s^E to \bar{p}_s^E for $s \geq \underline{s}^1$. This weakly increases the disposition to use in states 1 through $\underline{s}^1 - 1$ (Proposition 1, part (i-b)). Thus, the disposition to use in all states $s < \underline{s}^1$ is weakly higher with $\underline{\theta}$ than with $\bar{\theta}$. It follows that $(E, 1)$ continues to be an optimal choice in states $s < \underline{s}^1$ with $\bar{\theta}$, so the maximum first intentional use interval is weakly longer with $\bar{\theta}$ than with $\underline{\theta}$. Since nothing other than $(E, 1)$ is optimal in

states $s < \underline{s}_0^1$ with $\underline{\theta}$, nothing other than $(E, 1)$ can be optimal in states $s < \underline{s}_0^1$ with $\bar{\theta}$, so the minimum first intentional use interval is weakly longer with $\bar{\theta}$ than with $\underline{\theta}$. The argument is identical for a general decrease in p^A and for a general increase in u^A , b^A , or u^R .

Initial resistance interval. Suppose $\underline{p}_s^E \leq \bar{p}_s^E$ for all s (with all other parameters fixed). With $\underline{\theta}$, let \underline{s}^2 denote the first state (other than 0) in which $(R, 0)$ is not an optimal choice. Let $\underline{s}_0^2 \leq \underline{s}^2$ denote the first state (other than 0) in which something other than $(R, 0)$ is an optimal choice. Consider moving from $\underline{\theta}$ to $\bar{\theta}$ in two steps. (1) One state at a time, change from \underline{p}_s^E to \bar{p}_s^E for $s < \underline{s}^2$. Arguing as in the proof of Proposition 1, part (i-c), step 2, we see that when this change is made for state s , the value function is unchanged (so optimal choices are unchanged in all states other than s), $(R, 0)$ remains an optimal choice in state s , and the disposition to use weakly declines in state s . (2) Change from \underline{p}_s^E to \bar{p}_s^E for $s \geq \underline{s}^2$. This weakly reduces the disposition to use in state $s < \underline{s}^2$ (Proposition 1, part (i-a)). Thus, $(R, 0)$ continues to be an optimal choice in states $s < \underline{s}^2$ with $\bar{\theta}$, so the maximum initial resistance interval is weakly longer with $\bar{\theta}$ than with $\underline{\theta}$. Since nothing other than $(R, 0)$ is optimal in states $s < \underline{s}_0^2$ with $\underline{\theta}$, nothing other than $(R, 0)$ can be optimal in states $s < \underline{s}_0^2$ with $\bar{\theta}$, so the minimum initial resistance interval is weakly longer with $\bar{\theta}$ than with $\underline{\theta}$. The argument is identical for a general increase in p^A and for a general decrease in u^E , b^E , u^A , or b^A .

Final resignation interval. Suppose $\underline{p}_s^E \leq \bar{p}_s^E$ for all s (with all other parameters fixed). With $\underline{\theta}$, let \underline{s}^3 denote the first state (working backward from S) in which $(E, 1)$ is not an optimal choice. Let $\underline{s}_0^3 \geq \underline{s}^3$ denote the first state (working backward from S) in which something other than $(E, 1)$ is an optimal choice. Consider moving from $\underline{\theta}$ to $\bar{\theta}$ in two steps. (1) Change from \underline{p}_s^E to \bar{p}_s^E for $s > \underline{s}^3$. Since $(E, 1)$ is initially optimal for all such states, this leaves all optimal choices unchanged (Proposition 1, part (ii), coupled with the observation that, when $(E, 1)$ is optimal, neither $(A, 0)$ nor $(R, 0)$ is ever optimal). (2) Change from \underline{p}_s^E to \bar{p}_s^E for $s \leq \underline{s}^3$. This weakly increases the disposition to use in states $\underline{s}^3 + 1$ through S (Proposition 1, part (i-b)). Thus, the disposition to use in all states state $s > \underline{s}^3$ is weakly lower with $\underline{\theta}$ than with $\bar{\theta}$. It follows that $(E, 1)$ continues to be an optimal choice in states $s > \underline{s}^3$ with $\bar{\theta}$, so the maximum final resignation interval is weakly longer with $\bar{\theta}$ than with $\underline{\theta}$. Since nothing other than $(E, 1)$ is optimal in states $s > \underline{s}_0^3$ with $\underline{\theta}$, nothing other than $(E, 1)$ can be optimal in states $s > \underline{s}_0^3$ with $\bar{\theta}$, so the minimum final resignation interval is weakly longer with $\bar{\theta}$ than with $\underline{\theta}$. The argument is identical for a general increase in p^A and for a general decrease in u^A , b^A , or u^R . Q.E.D.

Proof of Proposition 3

Select any state s' . We can decompose the change from $\bar{\theta}$ to $\underline{\theta}$ into two components: (1) a change from $\bar{\theta}$ to θ' derived from $w_s(e, x, a) = \bar{w}_s(e, x, a) - d_{s'}$, and (2) a change from θ' to $\underline{\theta}$. The first change reduces u_s^a by $d_{s'}$ for all states s and actions a . This is simply a renormalization, and has no effect on choices. The second change weakly increases u_s^a by $d_{s'} - d_s$ for all $s < s'$, which weakly reduces the disposition to use in state s' by Proposition 1 part (i-b), and weakly decreases u_s^a by $d_s - d_{s'}$ for all $s > s'$, which also weakly reduces the disposition to use in state s' by Proposition 1 part (i-a). Thus, the disposition to use in state s' weakly decreases. Since this argument does not depend on the identity of s' , the disposition to use weakly decreases in all states. Q.E.D.

Proof of Proposition 4

Part (i). Consider some parameter vector $\bar{\theta}$, and let $\underline{\theta}$ denote the parameter vector obtained by setting $\underline{p}_s^a = 0$ for all a and s , leaving all other elements of $\bar{\theta}$ unchanged. By part (ii) of Proposition 1, continual use solves the DM's choice problem for $\bar{\theta}$ if and only if it does so for $\underline{\theta}$.

Part (ii). Consider some parameter vector $\bar{\theta}$, and suppose there is some state s' with $p_{s'}^E > 0$ such that $(E, 1)$ is not a best choice in s' . Applying (1) for $\theta = \bar{\theta}$ and using the fact that $(E, 1)$ is not a best choice in s' , we have

$$\begin{aligned} V_{s'}(\bar{\theta}) &= \max\{(1 - \bar{p}_{s'}^E)(\bar{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\bar{\theta})) \\ &\quad + \bar{p}_{s'}^E(\bar{u}_{s'}^E + \bar{b}_{s'}^E + \delta V_{\min\{S, s'+1\}}(\bar{\theta})) \\ &\quad (1 - \bar{p}_{s'}^A)(\bar{u}_{s'}^A + \delta V_{\max\{1, s'-1\}}(\bar{\theta})) + \bar{p}_{s'}^A(\bar{u}_{s'}^A + \bar{b}_{s'}^A + \delta V_{\min\{S, s'+1\}}(\bar{\theta})) \\ &\quad \bar{u}_j^R + \delta V_{\max\{1, s'-1\}}(\bar{\theta})\}. \end{aligned}$$

Since $(E, 1)$ is not a best choice in s' , (7) tells us that $\bar{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\bar{\theta}) > \bar{u}_{s'}^E + \bar{b}_{s'}^E + \delta V_{\min\{S, s'+1\}}(\bar{\theta})$, so the first term in braces is strictly less than $\bar{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\bar{\theta})$. By Assumption 2, the second term is strictly less than $(1 - \bar{p}_{s'}^A)(\bar{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\bar{\theta})) + \bar{p}_{s'}^A(\bar{u}_{s'}^E + \bar{b}_{s'}^E + \delta V_{\min\{S, s'+1\}}(\bar{\theta}))$, which is in turn less than $\bar{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\bar{\theta})$. Assumption 2 also implies that $\bar{u}_j^R + \delta V_{\max\{1, s'-1\}}(\bar{\theta})$ is strictly less than $\bar{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\bar{\theta})$. Thus, $V_{s'}(\bar{\theta}) < \bar{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\bar{\theta})$. Let $\underline{\theta}$ denote the parameter vector obtained by setting $\underline{p}_s^a = 0$ for all a and s , leaving all other parameters unchanged. Since the DM could select $(E, 0)$ in s' , we have $V_{s'}(\underline{\theta}) \geq \underline{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\underline{\theta}) \geq \bar{u}_{s'}^E + \delta V_{\max\{1, s'-1\}}(\bar{\theta})$, so $V_{s'}(\underline{\theta}) > V_{s'}(\bar{\theta})$. Q.E.D.

Proof of Proposition 5

First we argue that we can, without loss of generality, assume that the net transfer to each cohort is zero in each period. Consider an optimal policy $(\bar{\tau}, \bar{T})$ and an

associated optimal decision rule $\bar{\chi}$ that balances the government's budget. Define $\bar{L}_1^t = -\bar{\tau} \sum_{s=0}^S z_s^t(\bar{\chi}) \sigma_s^t(\bar{\chi})$ and $\bar{L}_2^t = T^t - \bar{L}_1^t$. The \bar{L}_1^t variables return the revenues raised from an age group to the same age group, while the variables \bar{L}_2^t are zero-sum inter-cohort transfers, with $\sum_{t=0}^{\infty} \pi^t \bar{L}_2^t = 0$. Now consider an alternative optimal tax problem that is identical to the original problem except that income in state s at age t is given by $\bar{y}_{st} = y_s - \bar{L}_2^t$. We claim that $(\bar{\tau}, \bar{L}_1^t)$ is a solution to this problem. For if some other feasible policy $(\tilde{\tau}, \tilde{L}_1^t)$ is superior, then, for the original problem, $(\tilde{\tau}, \tilde{L}_1^t + \bar{L}_2^t)$ is both feasible and superior to $(\bar{\tau}, \bar{L}_1^t)$, a contradiction. Notice that, for the alternative problem, $(\bar{\tau}, \bar{L}_1^t)$ belongs to the class of policies (τ, T) for which the net transfer to each cohort is zero in each period (i.e. that achieve budget balance both within age groups and within periods):

$$\tau \sum_{s=0}^S z_s^t(\chi) \sigma_s^t(\chi) = -T^t \quad (11)$$

for all t , where χ is optimal and balances the budget given (τ, T) .

Henceforth, we proceed as if $\bar{L}_2^t = 0$, implicitly incorporating the optimal inter-cohort transfers into income.

We begin with a lemma.

Lemma 3: *For all τ there exists $T \equiv (T^0, T^1, \dots) \in [-|\tau|, |\tau|]^\infty$ such that (τ, T) is a feasible policy satisfying budget balance within both period and cohort (equation (11)).*

Proof: Since the (potentially randomized) choice in each age-state pair is an element of Δ^3 , an intertemporal decision rule is an element of $(\Delta^3)^\infty$. Fix the value of τ . Let $\Gamma : [-|\tau|, |\tau|]^\infty \rightarrow \mathbb{P} \subset R^\infty$ denote the function mapping each T to a parameter vector Θ . Let $\Upsilon : \mathbb{P} \rightrightarrows (\Delta^3)^\infty$ denote the correspondence mapping each Θ to optimal decision rules. Endowing $[-|\tau|, |\tau|]^\infty$, \mathbb{P} , and $(\Delta^3)^\infty$ with the product topology and applying standard arguments, one can show that Γ is continuous, and that Υ is convex valued and upper-hemicontinuous. Let $\Psi : (\Delta^3)^\infty \rightarrow [-|\tau|, |\tau|]^\infty$ denote the function mapping each decision rule to per capita revenues generated, at each age, from the addictive substance. This is a linear function. Let $\Omega : [-|\tau|, |\tau|]^\infty \rightrightarrows [-|\tau|, |\tau|]^\infty$ denote the composition of Γ , Υ , and Ψ ; from the preceding, we know that this is a convex-valued upper-hemicontinuous correspondence mapping each sequence of lump-sum taxes into sequences of per capita revenues generated by the tax on the addictive substance. Budget balance holds when $T + \Omega(T) = 0$. Thus, if T is a fixed point of the correspondence $-\Omega$, then (τ, T) is feasible. Since $-\Omega$ is convex-valued and upper-hemicontinuous, and since $[-|\tau|, |\tau|]^\infty$ endowed with the product topology is compact and convex, existence of a fixed point follows from the Kakutani Fixed Point Theorem. Q.E.D.

Now we prove the proposition.

Part (i). We prove this in two steps.

Step 1. Suppose there is an optimal policy with $\tau > 0$ and that, contrary to the proposition, Condition A is satisfied for the decision rules leading to balanced budgets. We will establish a contradiction by showing that this policy must be strictly inferior to ϕ .

Let χ be an optimal decision rule satisfying budget balance for the original policy. We will show that χ yields a strictly higher age 0 discounted expected payoff with ϕ . Thus, the *optimal* choice with ϕ necessarily achieves a strictly higher discounted expected payoff than the optimal choice with the original policy.

Since the DM definitely starts at state 0 at age 0, and since there is some subsequent age at which at least two addictive states are reached with positive probability, there must be at least one intervening age at which neither use nor non-use is a certainty from the perspective of age 0.

Define $b^t \equiv \sum_{s=1}^S z_s^t(\chi) \sigma_s^t(\chi)$ (the probability of use at age t from the perspective of age 0). Consider any age t' at which neither use nor non-use is a certainty from the perspective of period 0 (that is, $b^{t'} \in (0, 1)$). We compute the expectation, as of age zero, of the marginal utility of income in s , t' , conditional on use and non-use of the addictive substance, under the policy ϕ , but assuming that the DM nevertheless continues to follow χ :

$$E_0[u'(e^{t'}) \mid x^{t'} = 0] = \frac{\sum_{s=1}^S z_s^{t'}(\chi)(1 - \sigma_s^{t'}(\chi))u'(y_s)}{1 - b^{t'}}$$

and

$$E_0[u'(e^{t'}) \mid x^{t'} = 1] = \frac{\sum_{s=1}^S z_s^{t'}(\chi)\sigma_s^{t'}(\chi)u'(y_s - q)}{b^{t'}} > \frac{\sum_{s=1}^S z_s^{t'}(\chi)\sigma_s^{t'}(\chi)u'(y_s)}{b^{t'}}$$

From this it follows that

$$E_0[u'(e^{t'}) \mid x^{t'} = 0] - E_0[u'(e^{t'}) \mid x^{t'} = 1] < \sum_{s=1}^S u'(y_s) z_s^{t'}(\chi) \left[\frac{1 - \sigma_s^{t'}(\chi)}{1 - b^{t'}} - \frac{\sigma_s^{t'}(\chi)}{b^{t'}} \right] \quad (12)$$

We know that

$$\sum_{s=1}^S z_s^{t'}(\chi) \left[\frac{1 - \sigma_s^{t'}(\chi)}{1 - b^{t'}} - \frac{\sigma_s^{t'}(\chi)}{b^{t'}} \right] = 0 \quad (13)$$

The term on the right hand side of (12) is simply a weighted sum of the terms in the summed in (13), where the weights increase weakly with s (since y_s weakly declines). We also know that $\frac{1 - \sigma_s^{t'}(\chi)}{1 - b^{t'}} - \frac{\sigma_s^{t'}(\chi)}{b^{t'}}$ is decreasing in the value of $\sigma_s^{t'}(\chi)$, and hence weakly decreasing in s (by hypothesis). Since the summation in (13) equals zero, the weighted sum in (12) must therefore be no greater than zero (it weakly shifts relative weight from every positive term to every negative term). It therefore follows that $E_0[u'(e^{t'}) \mid x^{t'} = 1] > E_0[u'(e^{t'}) \mid x^{t'} = 0]$.

Suppose we switch to ϕ . Assume for the moment the DM continues to follow χ . From the perspective of age 0, the result is an actuarially fair redistribution across age t realizations of (s, ω) . The DM receives the amount $\tau(1 - b^t) > 0$ in all (s, ω) for which $x_t = 1$, and gives up the amount τb^t in all (s, ω) for which $x_t = 0$ is assigned. If $b^t \in \{0, 1\}$, there is no redistribution and no effect on discounted expected hedonic payoff for t . When $b^t \in (0, 1)$, since $E_0[u'(e^t) \mid x^t = 1] > E_0[u'(e^t) \mid x^t = 0]$ for the last dollar redistributed, and since u is strictly concave, the transfer makes him strictly better off. Under our hypotheses, his discounted expected hedonic payoff weakly increases for every age t and strictly increases for some. Reoptimizing the decision rule reinforces this conclusion.

Step 2. Suppose ϕ is an optimal policy and that, contrary to the proposition, Condition A is satisfied for all optimal decision rules consistent with budget balance (with ϕ , this is all optimal decision rules). We establish a contradiction by showing that there are policies with strictly negative tax rates that are strictly superior to ϕ .

Choose some $\kappa \in (0, y_S - q)$. Consider some function $\hat{T} : [-\kappa, \kappa] \rightarrow [-\kappa, \kappa]^\infty$ such that $(\tau, \hat{T}(\tau))$ is a feasible policy with an optimal decision rule satisfying (11) for each τ (the existence of which is guaranteed by Lemma 3), and let Θ_τ be the parameter vector corresponding to this policy. Note that, as $\tau \rightarrow 0$, $\Theta_\tau \rightarrow \Theta^0$, where Θ^0 denotes the parameter vector corresponding to the policy ϕ .

Consider some sequence of tax rates and decision rules (τ_j, χ_j) with $\tau_j \rightarrow 0$, $\tau_j \in (-\kappa, 0)$, and $\chi_j \in \Upsilon(\Theta_{\tau_j})$ where χ_j satisfies (11) for $(\tau_j, \hat{T}(\tau_j))$, and where χ_j converges to a limit χ_∞ (this is possible since $(\Delta^3)^\infty$ endowed with the product topology is compact). Since Υ is upper-hemicontinuous, $\chi_\infty \in \Upsilon(\Theta^0)$.

Let b_j^t denote the likelihood of use at age t with χ_j (that is, $b_j^t \equiv \sum_{s=0}^S z_s^t(\chi_j) \sigma_s^t(\chi_j)$). Plainly, b_j^t converges to b_∞^t . From the government budget constraint, we know that $b_j^t = -\frac{\hat{T}^t(\tau_j)}{\tau_j}$ ($j \neq \infty$).

Choose some scalar b and consider the effect on expected hedonic payoffs of imposing a small substance tax τ coupled with a lump-sum tax $-b\tau$ at age t , starting from the parameter vector Θ^0 , and holding the decision rule fixed at χ_∞ . (This need not balance the government's budget.) Let $\zeta^t(\tau, \pi)$ denote the resulting age t expected instantaneous hedonic payoff (evaluated as of age 0). If $b_\infty^t \in (0, 1)$, then

$$\zeta^t(\tau, b) = (1 - b_\infty^t) E_0[u(e^t + b\tau) \mid x^t = 0] + b_\infty^t E_0[u(e^t + b\tau - \tau) \mid x^t = 1]$$

(where, in taking expectations, e^t is treated as the random variable, and its distribution corresponds to that implied by (ϕ, χ_∞)). If $b_\infty^t = 0$, then $\zeta^t(\tau, b) = E_0[u(e^t + b\tau)]$, and if $b_\infty^t = 1$, then $\zeta^t(\tau, b) = E_0[u(e^t + b\tau - \tau)]$. In the following, we use subscripts of ζ^t to denote partial derivatives.

We claim that $\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j}$ converges to $-\zeta_1^t(0, b_\infty^t)$. In other words, to a first order approximation, we can evaluate the age t welfare effect of switching from ϕ to

$(\tau_j, \widehat{T}(\tau_j))$ with the decision rule fixed at χ_∞ by computing the age t welfare effect of switching from ϕ to $(\tau_j, (-b_\infty^0 \tau_j, -b_\infty^1 \tau_j, \dots))$ with the decision rule fixed at χ_∞ . Obviously, $\frac{\zeta^t(\tau_j, b_\infty^t) - \zeta^t(0, b_\infty^t)}{\tau_j}$ converges to $-\zeta_1^t(0, b_\infty^t)$ (recall that $\tau_j < 0$), so the claim follows as long as $\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j}$ converges to $\frac{\zeta^t(\tau_j, b_\infty^t) - \zeta^t(0, b_\infty^t)}{\tau_j}$. Note that

$$\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j} - \frac{\zeta^t(\tau_j, b_\infty^t) - \zeta^t(0, b_\infty^t)}{\tau_j} = \int_{b_\infty^t}^{b_j^t} \left(\frac{1}{\tau_j} \int_0^{\tau_j} \zeta_{1,2}^t(\tau, b) d\tau \right) db$$

Since the second derivative of u is bounded, the term in parentheses remains finite as $j \rightarrow \infty$, so the left-hand side converges to zero.

Now we evaluate $\zeta_1^t(0, b_\infty^t)$ by taking the derivative with respect to τ and substituting $\tau = 0$ and $b = b_\infty^t = \sum_{s=0}^S z_s^t(\chi_\infty) \sigma_s^t(\chi_\infty)$. When $b_\infty^t \in (0, 1)$, we obtain

$$\begin{aligned} \zeta_1^t(0, b_\infty^t) &= (1 - b_\infty^t) b_\infty^t (E_0[u'(e^t) \mid x^t = 0] - E_0[u'(e^t) \mid x^t = 1]) \\ &< 0 \end{aligned}$$

(where we establish the inequality precisely as in step 1). For either $b_\infty^t \in \{0, 1\}$, we obtain $\zeta_1^t(0, b_\infty^t) = 0$.

Now suppose that the claim is false. Then, by step 1, ϕ is an optimal policy, and by hypothesis there is some age t' such that $b_\infty^{t'} \in (0, 1)$ with (χ_∞, Θ^0) . We will evaluate the payoff consequences of switching from (χ_∞, Θ^0) (policy ϕ along with the optimal decision rule χ_∞) to $(\chi_j, \Theta_{\tau_j})$ (policy $(\tau_j, \widehat{T}(\tau_j))$ along with the optimal decision rule χ_j) in two steps. First, switch from (χ_∞, Θ^0) to $(\chi_\infty, \Theta_{\tau_j})$ (that is, we change the policy without changing the decision rule). Second, switch from $(\chi_\infty, \Theta_{\tau_j})$ to $(\chi_j, \Theta_{\tau_j})$ (that is, we change the decision rule).

Consider the first switch, from (χ_∞, Θ^0) to $(\chi_\infty, \Theta_{\tau_j})$. We claim that, for j sufficiently large, this strictly increases discounted expected hedonic payoff. We divide the set of ages into two mutually exclusive and exhaustive subsets: those with $b_\infty^t \in \{0, 1\}$, and those with $b_\infty^t \in (0, 1)$. For ages in the first set, $\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j}$ converges to zero (since, as shown above, $\zeta_1^t(0, b_\infty^t) = 0$). For ages in the second set, $\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j}$ converges to $-\zeta_1^t(0, b_\infty^t) > 0$ (again, as shown above). Therefore, (a) for any finite $t^* \geq t'$, $\sum_{t=0}^{t^*} \delta^t \left(\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j} \right)$ converges to a number no smaller than $-\delta^{t'} \zeta_1^{t'}(0, b_\infty^{t'}) > 0$. Since $-u'(y_S - q - \kappa)$ is a lower bound on $\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j}$ (the policy can do no worse than reduce income by τ_j units in all states), we know that (b) $\sum_{t > t^*} \delta^t \left(\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j} \right) \geq -\frac{\delta^{t^*}}{1-\delta} u'(y_S - q - \kappa)$. Choose t^* such that (c) $-\frac{\delta^{t'} \zeta_1^{t'}(0, b_\infty^{t'})}{2} - \frac{\delta^{t^*}}{1-\delta} u'(y_S - q - \kappa) > 0$. By (a), we can then choose j^* such that (d) $\sum_{t=0}^{t^*} \delta^t \left(\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j} \right) > -\frac{\delta^{t'} \zeta_1^{t'}(0, b_\infty^{t'})}{2}$ for $j > j^*$. Summing (b) and (d) and using (c), we have $\sum_{t=0}^\infty \delta^t \left(\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j} \right) > 0$ for $j > j^*$. This implies that, for j sufficiently large, the switch from (χ_∞, Θ^0) to $(\chi_\infty, \Theta_{\tau_j})$ strictly increases discounted expected hedonic payoff, as claimed.

Next consider the switch from $(\chi_\infty, \Theta_{\tau_j})$ to $(\chi_j, \Theta_{\tau_j})$. The impact on age 0 expected discounted hedonic payoff is plainly bounded below by zero.

Putting these changes together, we see that the impact on age 0 expected discounted hedonic payoff is strictly positive for j sufficiently large, which contradicts the supposition that ϕ is an optimal policy.

Part (ii). We prove this in two steps, parallel to part (i).

Step 1. Suppose there is an optimal policy with $\tau < 0$ and that, contrary to the proposition, Condition B is satisfied for the decision rules leading to balanced budgets. We will establish a contradiction by showing that this policy must be strictly inferior to ϕ .

Let χ be an optimal decision rule satisfying budget balance for the original policy. We will show that χ yields a strictly higher age 0 discounted expected payoff with ϕ . Thus, the *optimal* choice with ϕ necessarily achieves a strictly higher discounted expected payoff than the optimal choice with the original policy. Let $\alpha_s^t(\chi)$ denote the probability of choosing $(R, 0)$ in (s, t) while following decision rule χ .

Consider any age t' for which $b^{t'} \in (0, 1)$. We compute the expectation, as of age zero, of the marginal utility of income in s, t' , conditional on use and non-use of the addictive substance, under the policy ϕ , but assuming that the DM nevertheless continues to follow χ :

$$\begin{aligned} E_0[u'(e^{t'}) \mid x^{t'} = 0] &= \frac{\sum_{s=1}^S z_s^{t'}(\chi) \left(\alpha_s^{t'}(\chi) u'(y_s - r_s) + (1 - \sigma_s^{t'}(\chi) - \alpha_s^{t'}(\chi)) u'(y_s) \right)}{1 - b^{t'}} \\ &\geq \frac{\sum_{s=1}^S z_s^{t'}(\chi) (1 - \sigma_s^{t'}(\chi)) u'(y_s)}{1 - b^{t'}} \end{aligned}$$

and

$$E_0[u'(e^{t'}) \mid x^{t'} = 1] = \frac{\sum_{s=1}^S z_s^{t'}(\chi) \sigma_s^{t'}(\chi) u'(y_s - q)}{b^{t'}}$$

(where we have used $q = 0$). Thus,

$$\begin{aligned} E_0[u'(e^{t'}) \mid x^{t'} = 0] - E_0[u'(e^{t'}) \mid x^{t'} = 1] &\geq \quad (14) \\ \sum_{s=1}^S u'(y_s) z_s^{t'}(\chi) \left[\frac{1 - \sigma_s^{t'}(\chi)}{1 - b^{t'}} - \frac{\sigma_s^{t'}(\chi)}{b^{t'}} \right] &+ \sum_{s=1}^S (u'(y_s) - u'(y_s - q)) z_s^{t'}(\chi) \frac{\sigma_s^{t'}(\chi)}{b^{t'}} \end{aligned}$$

We know that $\frac{1 - \sigma_s^{t'}(\chi)}{1 - b^{t'}} - \frac{\sigma_s^{t'}(\chi)}{b^{t'}}$ is decreasing in the value of $\sigma_s^{t'}(\chi)$, and hence weakly increasing in s (by hypothesis). Since the summation in (13) equals zero, the weighted sum comprising the first term on the right-hand side of (14) must be strictly positive (with y_s and $\sigma_s^{t'}(\chi)$ weakly decreasing and non-constant over states reached with positive probability, the weight on every strictly positive term weakly increases – and in some cases strictly increases – relative to the weight on every strictly negative term). It therefore follows that $E_0[u'(e^{t'}) \mid x^{t'} = 0] > E_0[u'(e^{t'}) \mid x^{t'} = 1]$ for q sufficiently small (the second term vanishes since u' is continuously differentiable).

Suppose we switch to ϕ . Assume for the moment the DM continues to follow χ . Arguing exactly as in step 1 of part (i), we see that the DM's discounted expected hedonic payoff weakly increases for every age t and strictly increases for some. Reoptimizing the decision rule reinforces this conclusion.

Step 2. Suppose ϕ is an optimal policy and that, contrary to the proposition, Condition B is satisfied for all optimal decision rules consistent with budget balance (with ϕ , this is all optimal decision rules). We establish a contradiction by showing that, if q is sufficiently small, there are policies with strictly positive tax rates that are strictly superior to ϕ .

Consider a sequence of tax rates and decision rules (τ_j, χ_j) with the same properties as in step 2 of part (i) except that $\tau_j \in (0, \kappa)$. Arguing as before, we see that: $\frac{\zeta^t(\tau_j, b_j^t) - \zeta^t(0, b_j^t)}{\tau_j}$ converges to $\zeta_1^t(0, b_\infty^t)$ as $j \rightarrow \infty$; $\zeta_1^t(0, b_\infty^t) = 0$ if $b_\infty^t \in \{0, 1\}$; and

$$\zeta_1^t(0, b_\infty^t) = (1 - b_\infty^t) b_\infty^t (E_0[u'(e^t) \mid x^t = 0] - E_0[u'(e^t) \mid x^t = 1])$$

if $b_\infty^t \in (0, 1)$. Using the same arguments as in step 1, we see that $\zeta_1^t(0, b_\infty^t) > 0$ for $b_\infty^t \in (0, 1)$ provided q is sufficiently small.

Now suppose that the claim is false. Then, by step 1, ϕ is an optimal policy, and by hypothesis there is some age t' such that more than one state is reached with positive probability, and neither y_s nor $\sigma_s^{t'}(\chi_\infty)$ is constant over these states (which in turn implies $b_\infty^{t'} \in (0, 1)$). As in step 2 of part (i), we evaluate the payoff consequences of switching from (χ_∞, Θ^0) to $(\chi_j, \Theta_{\tau_j})$ in two steps. For completely parallel reasons, a switch from (χ_∞, Θ^0) to $(\chi_\infty, \Theta_{\tau_j})$ is strictly beneficial for j sufficiently large, and a switch from $(\chi_\infty, \Theta_{\tau_j})$ to $(\chi_j, \Theta_{\tau_j})$ is weakly beneficial, so the two changes combined are strictly beneficial. This contradicts the supposition that ϕ is an optimal policy. Q.E.D.

Proof of Proposition 6

Arguing analogously to lemma 3, one can show that for all β there exists $T \equiv (T^0, T^1, \dots) \in [-|\beta|, |\beta|]^\infty$ such that (β, T) is a feasible policy satisfying budget balance both within period and with cohort.

We prove the result in two steps.

Step 1: A small steady-state rehabilitation subsidy (without net inter-cohort transfers) is beneficial.

Choose some $\kappa \in (0, y_S - q)$. Consider some function $\tilde{T} : [-\kappa, 0] \rightarrow [0, \kappa]^\infty$ such that $(\beta, \tilde{T}(\beta))$ is a feasible policy for each $\beta \in [-\kappa, 0]$, and let Θ_β be the parameter vector corresponding to this policy. Note that, as $\beta \rightarrow 0$, $\Theta_\beta \rightarrow \Theta^0$. Let $p_{s\beta}^a$ denote the value of p_s^a with policy β . Under Assumption 3, $p_{s\beta}^a \in [0, p_{s0}^a]$ for all s and $a = E, A$ (since the policy β weakly reduces net income). For any β , construct the parameter

vector Θ_β^0 by taking the probabilities of entering the hot mode (p_s^a) from Θ_β and all other parameters from Θ^0 .

Consider some sequence of tax rates and decision rules (β_j, χ_j) with $\beta_j \rightarrow 0$, $\beta_j \in (-\kappa, 0)$, and $\chi_j \in \Upsilon(\Theta_{\beta_j})$ where χ_j balances the budget for $(\beta_j, \tilde{T}(\beta_j))$. Without loss of generality, assume that the sequence of decision rules converges to a limit χ_∞ . Since Υ is upper-hemicontinuous, $\chi_\infty \in \Upsilon(\Theta^0)$. For notational convenience, define $\beta_\infty = 0$ (so that $\Theta_{\beta_\infty} = \Theta^0$).

Once again let $\alpha_s^t(\chi)$ denote the probability of choosing $(R, 0)$ in (s, t) while following decision rule χ . Let $z_{sj}^t(\chi)$ denote the probability (from the perspective of age 0) of reaching state s at age t following χ and using the hot mode probabilities $p_{s\beta_j}^a$. Let $B_{ij}^t = \sum_{s=0}^S z_{si}^t(\chi_j) \alpha_s^t(\chi_j)$; this is the probability of choosing rehabilitation at age t while adhering to χ_j under Θ_{β_j} , from the perspective of age 0. Plainly, B_{jj}^t and $B_{j\infty}^t$ both converge to $B_{\infty\infty}^t$. Note also that $B_{\infty\infty}^t \in (0, 1)$ iff $B_{j\infty}^t \in (0, 1)$, and $B_{\infty\infty}^t = 0$ iff $B_{j\infty}^t = 0$. From the government budget constraint, we know that $B_{jj}^t = -\frac{\tilde{T}^t(\beta_j)}{\beta_j}$ ($j \neq \infty$).

We decompose the effect of the switch from (χ_∞, Θ^0) to $(\chi_j, \Theta_{\beta_j})$ into three steps. First, we switch from (χ_∞, Θ^0) to $(\chi_\infty, \Theta_{\beta_j}^0)$ (that is, switch only the hot mode probabilities); second, switch from $(\chi_\infty, \Theta_{\beta_j}^0)$ to $(\chi_\infty, \Theta_{\beta_j})$ (that is, change the rest of the parameters); third, switch from $(\chi_\infty, \Theta_{\beta_j})$ to $(\chi_j, \Theta_{\beta_j})$ (that is, switch the decision rule).

The first step weakly reduces p_s^a for all s , $a = E, A$. We claim that this weakly increases discounted expected hedonic payoff. To prove this, imagine that the DM starts in some arbitrary state s , and consider changing these probabilities only for the current period, and not for any subsequent period. For any state s , entering the hot mode in period t matters only if χ_∞ prescribes either $(E, 0)$ or $(A, 0)$. In either case, the discounted expected payoff (from age t onward) received in the cold mode weakly exceeds the discounted expected payoff received in the hot mode (since χ_∞ is optimal given the original probabilities), so weakly reducing p_s^a for the current period weakly increases the discounted expected payoff. Thus, we can make the switch from Θ^0 to $\Theta_{\beta_j}^0$ by first changing probabilities for age 0, then for age 1, then for age 2, and so forth; since each step weakly increases the continuation payoff associated with every state, the collection of steps must weakly increase discounted expected payoff as of age 0.

Now consider the second step, from $\Theta_{\beta_j}^0$ to Θ_{β_j} . We show next that this strictly increases discounted expected hedonic payoff.

Let $\gamma^t(\beta, B, j)$ denote the age t instantaneous expected hedonic payoff (evaluated as of age 0) with policy (β, B) using the hot mode probabilities from Θ_{β_j} and holding the decision rule fixed at χ_∞ . That is, if $B_{j\infty}^t \in (0, 1)$ (equivalently $B_{\infty\infty}^t \in (0, 1)$),

$$\gamma^t(\beta, B, j) = (1 - B_{j\infty}^t) E_0^j[u(e^t + \beta B) \mid a^t \neq R] + B_{j\infty}^t E_0^j[u(e^t + \beta B - \beta) \mid a^t = R]$$

(where, in taking expectations, e^t is treated as the random variable, and its distribution corresponds to that implied by (χ_∞, ϕ) using the hot mode probabilities from Θ_{β_j}). If $B_{j\infty}^t = 0$ (equivalently $B_{\infty\infty}^t = 0$), then $\gamma^t(\beta, B, j) = E_0[u(e^t + \beta B)]$, and if $B_{j\infty}^t = 1$ (equivalently $B_{\infty\infty}^t = 1$), then $\gamma^t(\beta, B, j) = E_0[u(e^{t'} + \beta B - \beta)]$. In the following, we use subscripts of γ^t to denote partial derivatives.

Arguing as in step 2 of Proposition 5, part (i), one can show that $\frac{\gamma^t(\beta_j, B_{jj}^t, j) - \gamma^t(0, B_{jj}^t, j)}{\beta_j}$ converges to $-\gamma_1^t(0, B_{\infty\infty}^t, 0)$. For $B_{\infty\infty}^t \in \{0, 1\}$, we obtain $\gamma_1^t(0, B_{\infty\infty}^t, 0) = 0$. For $B_{\infty\infty}^t \in (0, 1)$, we obtain

$$\gamma_1^t(0, B_{\infty\infty}^t, 0) = B_{\infty\infty}^t (1 - B_{\infty\infty}^t) (E_0^\infty[u'(e^t) \mid a^t \neq R] - E_0[u'(e^t) \mid a^t = R]),$$

which we will show is strictly negative.

Under ϕ , the environment is stationary, so for each state s the set of best choices is the same for all t . By hypothesis, there is some s^* such that $(R, 0)$ is the unique best choice in s^* (so $\alpha_{s^*}^t(\chi_0) = 1$), $(R, 0)$ is not a best choice for $s < s^*$ (so $\alpha_s^t(\chi_0) = 0$), and $(E, 1)$ is not a best choice for some $s < s^*$. From this we reach two further conclusions: first, $z_{s\infty}^t(\chi_0) = 0$ for all (s, t) with $s > s^*$; second, there is some t' such that $z_{s^*\infty}^{t'}(\chi_0) \in (0, 1)$; third, for all t , $z_{s^*\infty}^t(\chi_0) > 0$ implies $B_{\infty\infty}^t = z_{s^*\infty}^t(\chi_0)$.

Next we show that $B_{\infty\infty}^t \in (0, 1)$ implies $E_0^\infty[u'(e^t) \mid a^t = R] > E_0^\infty[u'(e^t) \mid a^t \neq R]$, from which it follows immediately that $\gamma_1^t(0, B_{\infty\infty}^t, 0) < 0$, as claimed. Note that

$$E_0^\infty[u'(e^t) \mid a^t = R] = u'(y_{s^*} - r_{s^*})$$

and

$$E_0^\infty[u'(e^t) \mid a^t \neq R] = \frac{\sum_{s=1}^{s^*-1} z_{s\infty}^t(\chi) \left[(1 - \sigma_s^{t'}(\chi))u'(y_s) + \sigma_s^{t'}(\chi)u'(y_s - q) \right]}{\sum_{s=1}^{s^*-1} z_{s\infty}^t(\chi)}$$

Since $y_{s^*} \leq y_s$ for $s < s^*$ and $r_{s^*} > q$, we have $y_{s^*} - r_{s^*} < y_s - q < y_s$ for all $s < s^*$. The desired conclusion then follows directly from the fact that u is strictly concave.

We know that $B_{\infty\infty}^{t'} \in (0, 1)$. To complete the proof that the second step is strictly beneficial, one therefore uses an argument analogous to that appearing near the end of the proof of Proposition 5, part (i), step 2.

Finally, consider the switch from $(\chi_\infty, \Theta_{\beta_j})$ to $(\chi_j, \Theta_{\beta_j})$. The impact on age 0 expected discounted hedonic payoff is plainly bounded below by zero.

Putting these changes together, we see that the impact on age 0 expected discounted hedonic payoff is strictly positive for j sufficiently large.

Step 2: A small steady-state rehabilitation tax (without net inter-cohort transfers) is harmful.

For policy ϕ , let s^* denote the earliest state in which rehabilitation is an optimal choice, and let $s' \in \{1, \dots, s^* - 1\}$ denote a state in which $(E, 1)$ is not a best choice (both states are referenced in the proposition).

For any $\beta > 0$, let $\tilde{T}(\beta)$ be a feasible policy, and let χ_β be an associated optimal choice rule that achieves budget balance. Notice that, as $\beta \downarrow 0$, $(\theta_\beta^t, \theta_\beta^{t+1}, \dots)$ (the sequence of parameter vectors facing the DM starting at age t given the policy $(\beta, \tilde{T}(\beta))$) converges to Θ^0 uniformly over all t (since the absolute value of $\tilde{T}^t(\beta)$ cannot exceed β). Thus, there exists a threshold $\hat{\beta}$ such that when $\beta \in (0, \hat{\beta})$, for all t : $(R, 0)$ is the unique optimal choice in s^* , $(R, 0)$ is not an optimal choice in any $s < s^*$, and $(E, 1)$ is not an optimal choice in s' (this follows from the fact that Υ is upper-hemicontinuous, and the fact that if any mixture between two deterministic choices is optimal for some state, then all mixtures between these choices are also optimal). Accordingly, with the policy $(\beta, \tilde{T}(\beta))$ and the decision rule χ_β for $\beta \in (0, \hat{\beta})$, the DM never advances beyond state s^* , always rehabilitates in state s^* , never rehabilitates in any state $s < s^*$, and, for some t , reaches state s^* with some positive probability less than unity. Arguing as in step 1, we see that $E_0[u'(e^t) \mid a^t = R] < E_0[u'(e^t) \mid a^t \neq R]$, where we take expectations assuming the policy ϕ is in place, but the hot-mode probabilities associated with $(\beta, \tilde{T}(\beta))$ prevail and the DM continues to follow χ_β .

For $\beta \in (0, \hat{\beta})$, we evaluate the change from $(\beta, \tilde{T}(\beta))$ to ϕ in three steps. First, change the policy from (β_j, T_j) to ϕ without changing the hot mode probabilities, and keeping the choice rule fixed at χ_β . This creates an actuarially fair redistribution for each age t from realizations with no rehabilitation to realizations with rehabilitation. Since there is at least one t for which both rehabilitation and no rehabilitation are chosen with positive probability, this is strictly beneficial. Second, reoptimize the decision rule; this is weakly beneficial. Third, change the hot mode probabilities to those prevailing with the policy ϕ and reoptimize the decision rule. Since $\tilde{T}^t(\beta) \leq 0$ for all t , under Assumption 3 the hot mode probabilities weakly decline, so this is weakly beneficial. Q.E.D.