The desperation threshold: a model to explain decisions in poverty

Du Buisson de Courson, B.

Citation

Du Buisson de Courson, B. (2026, April 7). *The desperation threshold: a model to explain decisions in poverty*. Retrieved from https://hdl.handle.net/1887/4301048

3 Why does poverty increase time discounting? Present needs and uncertain future

3.1 Abstract

In situations of poverty, people are observed to heavily discount the future. Despite a clear empirical consensus, the theoretical literature lacks satisfying models explaining this social gradient. In particular, the urge of satisfying urgent needs is often invoked verbally as an explanation, but has never been properly modeled. Here, we show that the desperation threshold model can explain high discounting in situations of poverty. We present an analytical model with four different intertemporal scenarios of the desperation threshold, varying the consequences of the threshold on future utility. The model predicts high discounting around the threshold, and high patience at intermediate levels of resources. In contrast with the consequences of the desperation on risk taking, as the effect is predicted to occur on both sides of the threshold. We show that, unlike existing accounts, our explanation does not depend on assuming a future improvement, but predicts a U-shaped effect of future expectations. We discuss the relevance of our findings for social sciences and public policy.

3.2 Introduction

In situations of poverty, people are often observed to heavily discount the future (Carvalho, 2010; Green et al., 1996; Haushofer & Fehr, 2014; Lawrance, 1991; Reimers et al., 2009). A dramatic manifestation of this phenomenon is 'payday loans': each year, about ten million Americans with low incomes resort to short term borrowing, at an annual interest rate of the order of 400% (A. V. Banerjee & Duflo, 2007; Bertrand & Morse, 2011; Dobbie & Skiba, 2013). An association between poverty and time discounting has been observed repeatedly since Malthus (1798). Causality has been proven to be bidirectional (Doepke & Zilibotti, 2008): not only do short-term oriented individuals accumulate less wealth (Epper et al., 2020; Sunde et al., 2022), but low resources also increases time discounting (Carvalho et al., 2016; Handa et al., 2020; Haushofer et al., 2013) – the process we are here interested in. In a recent review, time discounting was found to be the only known causally robust effect of poverty on decision making (Haushofer & Salicath, 2023). This effect is of interest to several disciplines, including economics (Lawrance, 1991), sociology (Duvoux, 2023), psychology (Griskevicius et al., 2013; Haushofer & Fehr, 2014) and criminology (Gelder & Frankenhuis, 2025). It has profound implications: it has been proposed (i) as the common denominator of the 'behavioural constellation of poverty' - the set of behaviours typical of low-income populations (Pepper & Nettle, 2017), (ii) as a possible reason for poverty persistence (Haushofer & Fehr, 2014), and (iii) as a reason for the higher prevalence of crime - a behaviour that often brings short-term benefits and potential long-term legal troubles - in deprived populations (Brezina et al., 2009; Pepper & Nettle, 2017).

In stark contrast with the settled empirical case, it remains unclear why poverty increases time discounting. Two perspectives have been proposed. The first considers high discounting as a symptom that poverty degrades decision making. This can occur through stress (Haushofer & Fehr, 2014), ego-depletion (Spears, 2011) or, most famously, a 'tunneling effect': as Shah (Shah et al., 2012) put it, "While focusing on the groceries from week to week, we might neglect next month's rent" (p. 682). In a similar vein, A. V. Banerjee & Duflo (2007) suggest that not thinking about the future is "emotionally wise [...], to avoid confronting the sheer inadequacy of the standard of living" (p. 165).

A second perspective, which we adopt here, proposes that high discounting is not dysfunctional, but on the contrary an appropriate response to a context of poverty. A common argument concerns "collection risks": in poverty, people may be less likely to actually collect the future reward, for instance because of a higher mortality risk (Griskevicius et al., 2011). However, people with low incomes also exhibit high discount rates over very short periods of times (e.g. weeks in the case of payday loans), where mortality risks are negligible (Riis-Vestergaard & Haushofer, 2017). Furthermore, this class of explanations cannot account for the influence of 'states', like hunger (Allen & Nettle, 2021) or more generally financial need (Carvalho et al., 2016; Fitzpatrick & Coleman-Jensen, 2014; Sharma et al., 2023) on time discounting: collection risks are an environment-level parameter, presumably stable on short-term scales. In any case, collection risks are only one side of the problem: poverty also plausibly increases how much one needs resources right now, and can thereby generate 'waiting costs' (Mell et al., 2021). Many authors (Epper, 2015; Fisher, 1930; Frankenhuis & Nettle, 2020; Haushofer et al., 2013; Mell et al., 2021; Sharma et al., 2023) thus propose a simple verbal justification, well summarised by the character Earn in the TV series *Atlanta*: "poor people don't have time for investments, because poor people are too busy trying not to be poor! I need to eat today, not in September...".

However intuitive Earn's argument may be, we consider it underspecified and potentially misleading. The optimal level of time discounting depends not only on the value of present consumption, but also on the (expected) value of future consumption – which should also increase in situations of poverty, if low resources today tends to be correlated with low resources tomorrow. Simply put, Earn will also need to eat in September; his current level of need is therefore not a sufficient reason to discount the future. The reasoning only holds if survival is at stake (Chavas, 2013) or if Earn has reasons to expect his situation to improve in the meantime – that is, if his poverty is not a persistent condition, but a rough patch. We are aware of only two attempts to formalize this intuition, one in economics (Epper, 2015) and one in evolutionary psychology (Mell et al., 2021). Both rely on the assumption that individuals in poverty expect their situation to improve. If someone expects to be even worse off in the future, these models instead predict patience (Epper, 2015, p. 4). In reality, the "truly disadvantaged" populations (W. J. Wilson, 2012) tend to be stuck in poverty, and are often found by ethnographers to be rather pessimistic (Lewis, 1963; J. MacLeod, 2018). Among them, pessimism has been found to be associated with low saving (Gladstone & Pomerance, 2025), and in particular, offenders have been described in ethnographies as both extremely short term oriented and particularly pessimistic about the future (Anderson, 2000; Dickinson et al., 2025; W. J. Wilson, 2012). Thus, the literature lacks a model capable of explaining why people tend to discount the future when they struggle to meet basic needs, and their future is not expected to be brighter.

In this paper, we show that Earn's intuition is warranted under an alternative conceptualisation of needs', grounded in the desperation threshold model (DTM) (de Courson, Frankenhuis, Gelder, et al., 2025). The DTM posits that individuals experience a critical resource level representing basic needs above which they strongly prefer to be. In contrast with the standard microeconomics model, the DTM does not assume a concave mapping between resources and utility, but a S-shape. The utility function features a very steep region – the desperation threshold – where basic needs hang in the balance. It is flatter both above and below the threshold, as basic needs are respectively securely satisfied and definitively unmet. Here, we extend the DTM from the study of risk taking to the study of intertemporal decisions, introducing four distinct scenarios, varying the consequences of the desperation threshold on utility. We demonstrate that, regardless of implementation, one should discount the future more when close to this threshold. This occurs because the long term is more uncertain than the present: individuals situation may improve or worsen, both resulting in lower subjective resource valuation. For the same reason, three out of the four scenarios predicts even more pronounced discounting when a desperate individual has *directional* expectations about the future: whether she expects to have more or less resources in the future, future rewards become further devalued. Thus, unlike previous approaches (Epper, 2015; Mell et al., 2021), our result is not contingent on assuming that individuals in poverty are optimistic about the future, but it integrates those models highlighting optimism's role with evidence of extreme discounting among fatalistic populations (Anderson, 2000; Dickinson et al., 2025; M. Wilson & Daly, 1985).

3.3 Model

We model an agent living for two periods (Fig. 3.1A), defined by a single state parameter: its level of resources, that changes stochastically over time. The agent begins with x_0 resources. The resource level is subject to random shocks before the first time period, and between the first and the second. These shocks are independent and identically distributed, following a normal distribution with mean zero and variance σ^2, with σ representing environmental volatility. Between periods, a constant drift parameter d is applied to the resource level. This parameter captures future expectations: $\mathbb{E}(x_2|x_1) = x_1 + d$, in other words one expects his state to improve over time if $d > 0$, and to deteriorate if $d < 0$. For simplicity, we do not distinguish between consumption, income and wealth: the variables x_i summarises the material standing of the agent at time i.

The agent's preferences are represented by a utility function that depends on resource states in both periods, x_1 and x_2. Here, we realized that the desperation threshold concept could be translated into an intertemporal utility function in multiple ways. The threshold is verbally defined as a strong disutility below some level of resources (de Courson, Frankenhuis, Gelder, et al., 2025). Yet, it has been formalised in several, slightly different ways (de Courson et al., 2023; de Courson, Frankenhuis, & Nettle, 2025; de Courson & Nettle, 2021). While these implementations yield qualitatively similar predictions regarding risk taking — the primary dependent variable in previous studies — their implications on time discounting may differ substantially. We identified four subtly different scenarios of the DTM idea within our two-period setup. For the sake of exhaustiveness, we present all four scenarios, highlighting the places where the results diverge. We keep these functions as simple as possible, combining a linear dependence on states with thresholds effects. In other words, we neglect diminishing marginal utility at high resource levels, since our analysis focuses on decision making around the threshold.

In Fig. 3.1, we define the four utility functions, along with conceptual figures and their intuitive names. In the first one, falling below the threshold incurs a fixed utility penalty w. This penalty applies additively: being below the threshold in both periods generates twice the disutility of being below once. Importantly, utility continues to decrease linearly below the threshold without a lower bound, meaning that one is still sensitive to its level of resources even if can not meet the goal. The second function sets period utility to zero whenever resources fall below the threshold. Here and in the next functions, an agent becomes indifferent to its state when below the threshold: if one is below, it does not matter how far below one is. To create a discontinuous utility jump at the threshold, we introduce a constant also denoted w in this and subsequent functions. This function can represent 'hibernation': being below the threshold makes it impossible to accumulate utility in the present, but preserves the possibility to do so in the future. The third function represents a 'game-over' below the threshold: it cancels not only the utility of this time period, but also of possible future time periods. This captures irreversible consequences, such as eviction: even if resources become sufficient tomorrow, today's eviction cannot be undone. Our fourth function sets lifetime utility to zero if resources fall below the threshold in any period. Even if one has been above the threshold in the past, it has all been 'in vain'. This can for instance capture starvation during development: life terminates, and the previously accumulated physical capital is inconsequential.

These utility functions give rise to a time orientation, i.e. a preference for sooner, or later, resources. We capture the agent's level of time discounting through its marginal rate of intertemporal substitution (MRIS) – the amount one must receive to willingly surrender one resource unit today. Formally,

$$MRIS(x_0) = \frac{\frac{d\mathbb{E}(U)}{dx_1}}{\frac{d\mathbb{E}(U)}{dx_2}}$$

In other words, the agent decides whether to discount the future based on her initial level of resources, by comparing to what extent an extra resource in the first time period makes her happier than an extra resource in the second one. Note that economists often distinguish between genuine time discounting – an exogenous disregard for future consequences just because they happen in the future – from rational motives to prefer immediate rewards (Carvalho et al., 2016; Haushofer & Salicath, 2023). Here, we do not make such a distinction, as we aim to find rational explanations for observed behaviour. Note also that when computing the marginal utility in the first period, we assume that the extra resource is not present in the second period. This is simply a way to present the results in a more insightful way. Otherwise, $MRIS > 1$ by construction: one always prefers to have the resource in the short term, even if it is mainly useful in the long term. Note also that x_1 and x_2 are random variables. $MRIS$ thus depends on *expected* marginal utilities knowing only x_0. By differentiating with respect to x_1, we mean with respect to its expectation, while keeping $\mathbb{E}(x_2)$ constant. This specification is necessary to obtain a meaningful discount factor: as we study discontinuous – and therefore non-differentiable – utility functions, the random perturbations smoothes utility, ensuring marginal utility is well-defined.

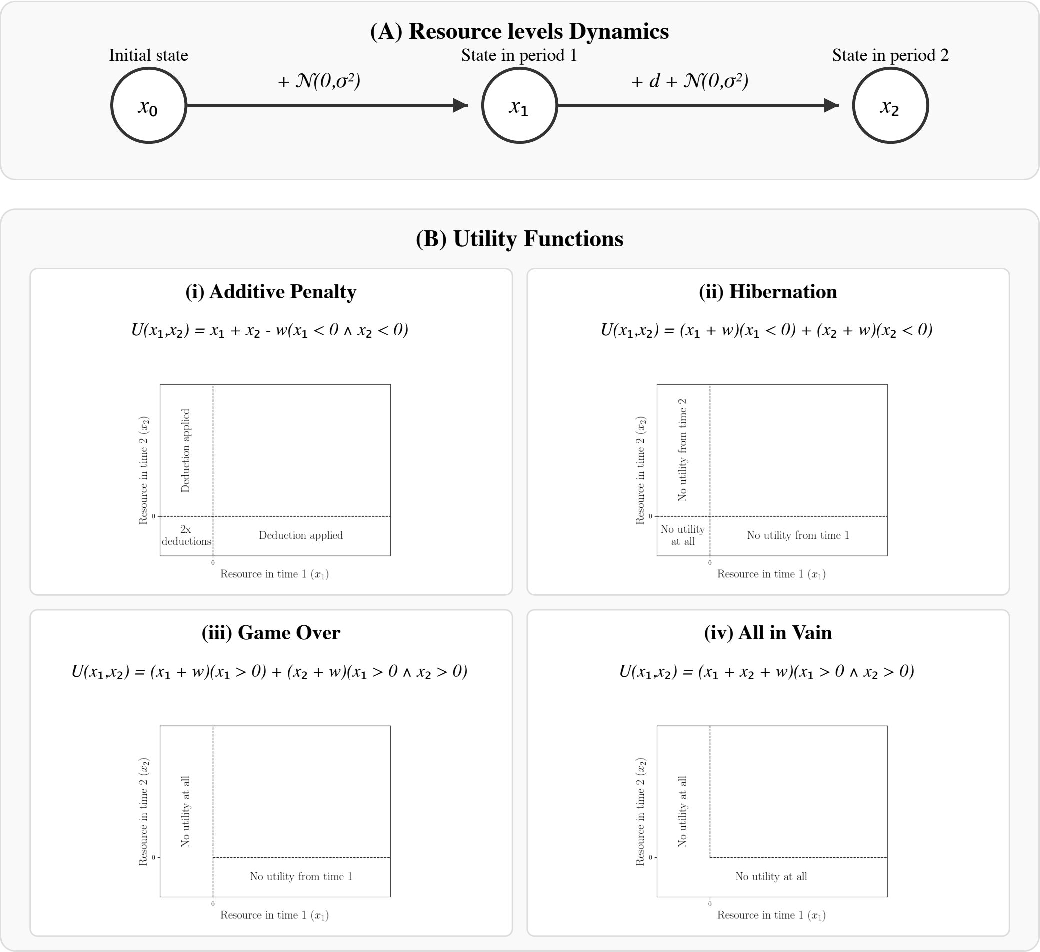

Figure 3.1: Summary of the model assumptions. Utility functions are visualized in heatmaps, where lighter tones indicate higher utility. Zero marks the desperation threshold, hatches mark areas where utility is zero.

3.4 Results

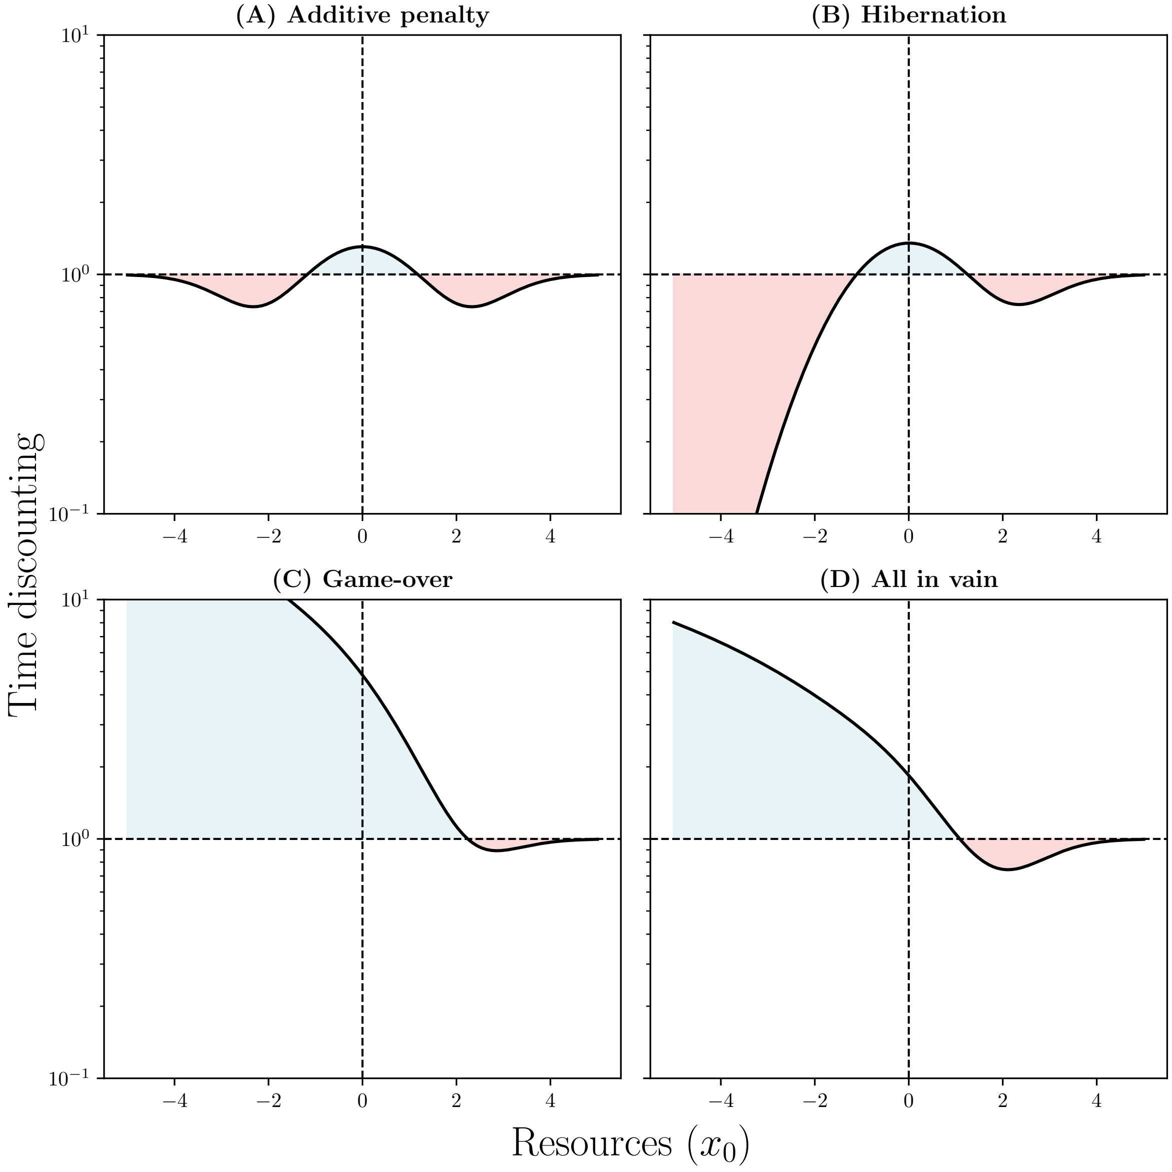

Figure 3.2: Time discounting depending on the initial level of resources x_0. Here, we assume neutral future expectations ($d = 0$), and $\sigma = 1$. The vertical dashed line marks the desperation threshold, the horizontal one marks time neutrality. The y-axis is truncated at .1 and 10 for legibility.

3.4.1 Analytical derivation

We analyse our model by describing how time discounting, as measured by the MRIS varies depending on the agent's state x_0, future expectations d and environment volatility σ. Here, we give an intuitive account of the results, and develop some mathematical details of only the simpler cases – the 'additive penalty' and the 'hibernation' scenario, that are time-separable. In the appendix, we present detailed analytical derivations for all four scenarios, asymptotic approximations and numerical sensitivity analysis.

To derive marginal utilities formulas, we note that an extra resource has two effects on the agent's utility. First, a 'level effect': it makes her richer on average. Second, the extra resource has a 'threshold effect': it may push one above the threshold, which brings a utility of w with a weight $\phi(x_0)$ – the normal probability density function. This component follows the bell-shaped Gaussian curve, with a peak at $x_0 = 0$. This captures the DTM intuition: one never needs resources more than on the threshold of desperation. The marginal utility decomposes naturally using the Leibniz rule: the 'level effect' comes from differentiating under the integral, while the 'threshold effect' arises from the change in probability mass at the threshold – a boundary term. In the second time period, the agent faces a similar dilemma, hence the same functional form, only with a doubled variance due to the longer time horizon, and a drift d. In the additive penalty case, the MRIS simplifies to:

$$MRIS = \frac{1 + w\phi_{\sigma^2}(x_0)}{\underbrace{1}_{\text{level effect}} + \underbrace{w\phi_{2\sigma^2}(x_0 + d)}_{\text{threshold effect}}}$$

3.4.2 The effect of resources

In Fig. 3.2, we plot discounting levels depending on x_0, assuming neutral future expectations ($d = 0$). Above the threshold, all four utility scenarios yield qualitatively similar patterns. Far above the threshold, the agent exhibits time neutrality ($MRIS \sim 1$). This occurs because the threshold becomes irrelevant ($\Phi(-x_0) = o(1)$), and by assumption, resources in both periods yield equal utility. Slightly above the threshold, the agent should be patient ($MRIS < 1$), out of precaution: immediate needs are almost certainly met, but future needs remain at stake.

The agent becomes impatient ($MRIS > 1$) when x_0 is very close to the threshold. In this region, meeting the threshold hangs in the balance both periods, but greater uncertainty about the future makes immediate gains more important. More precisely, one can thus be quite sure to need the resources in the short term, but in the long term, her situation is likely to have changed more. Whether her future condition improves or worsens, the probability of being near the threshold declines (Fig. 3.3), reducing the marginal value of future resources.

Fig. 3.3 illustrates these two results by comparing the probability that a marginal resource shifts the agent above the threshold in each period. Analytically, for $x_0 > 0$, $\phi_{2\sigma^2}(x_0) > \phi_{\sigma^2}(x_0) \iff x_0 < \sigma\sqrt{2\ln 2}$. This defines the region in which the agent prefers present resources in the 'additive penalty' case: $MRIS > 1 \Leftrightarrow x_0 < \sigma\sqrt{2\ln 2}$ (Fig. 3.3). In the 'hibernation' scenario, it is a sufficient condition for $MRIS > 1$ (see SI). This implies that size of this 'impatience' region is roughly proportional to the volatility σ: higher uncertainty requires a higher buffer to feel secure in the present. However, volatility does not affect overall time preference directionally. Instead, it acts like a scaling parameter, expanding both precautionary and impatient zones (Fig. 3.5). The pattern is also robust to changes in the parameter w (Fig. 3.6).

These results are related to the more general concept of 'prudence', defined in utility theory as the convexity of marginal utility – formally, the third derivative of the utility function, usually assumed to be positive. Under this assumption, uncertainty increases expected marginal utility: by Jensen's inequality, if U' is convex $\mathbb{E}(U'(x)) \geq U'(\mathbb{E}(x))$. This leads to 'precautionary savings': when future resources are on average equal to present ones but uncertain, one should save, as the risk of being worse off tomorrow outweighs the benefits of being better off. However, in our model, marginal utility is bell-shaped in all four scenarios, and thus convex for high values – hence patience at those level of resources – but concave around the desperation threshold, at the cusp of the bell. This point of view allows to recover the above result in the time-separable scenarios: in the 'additive penalty' scenario, $U'''(x) > 0 \iff |x| < \sigma$. Like in our model, this results in time discounting in a region centered around the threshold, and of width proportional to σ, but $\sqrt{2\ln 2}$ smaller – as this compares marginal utility in time 1 and $1+\epsilon$ rather than time 1 and time 2 as in the above model. A concave marginal utility reverts the Jensen inequality: $\mathbb{E}(U'(x)) \leq U'(\mathbb{E}(x))$, meaning that uncertainty makes resources less useful in the future, in expectation. This reasoning is analogous to the one-dimensional heat equation: in our framework, time has a 'diffusive' effect, it brings randomness, and therefore spreads marginal utility (Fig. 3.3). Like in the heat equation ($\frac{df}{dt} \propto \frac{d^2f}{d^2x}$), a resource is expected to be less useful in the future if and only if expected marginal utility is concave at that point.

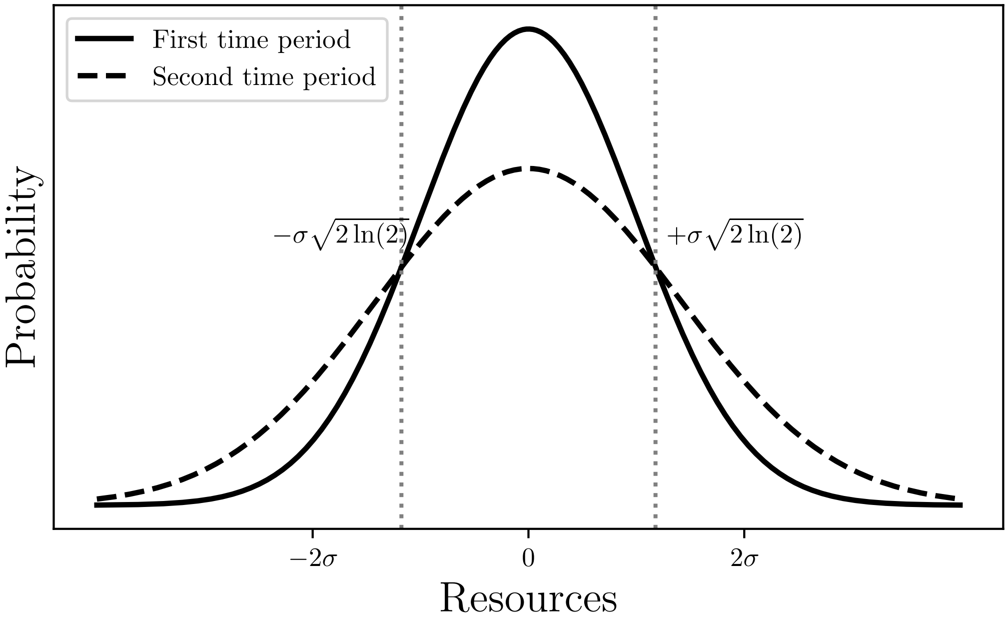

Figure 3.3: Normal probability density functions ϕ_{σ^2} and $\phi_{2\sigma^2}$. They represent the probability density function of x_1 and x_2, assuming $x_0 = 0$ and $d = 0$. Equivalently, due to the Gaussian density symmetry, they describe the likelihood that future resource levels x_1 and x_2 are exactly zero, given $x_0 = x$. This yields the 'threshold effect' (see Results) on marginal utility: the probability that an extra resource pushes one above the threshold. In the central region ($|x| < \sqrt{2ln(2)}$), an extra resource is more valuable to cross the threshold in the first period than in the second, and vice versa outside this region.

Far below the threshold, the four utility scenarios diverge in their predictions. The 'additive penalty' scenario predicts a symmetric pattern, in line with the symmetry of its utility function. When the threshold is out of reach, either because one is definitely above it, or definitely below it, utility becomes linear in this scenario (Fig. 3.1A), and the agent acts as if the threshold did not exist. In the 'hibernation' case, the agent becomes infinitely patient: she receives no utility today, but keeps some hope for recovery tomorrow. Formally, since uncertainty increases over time, the agent has a higher chance of surpassing the threshold in the second period than in the first: as $x_0 \to -\infty$, $\phi_{\sigma^2}(x_0) = o(\phi_{2\sigma^2}(x_0))$ (Fig. 3.3). The best strategy is to accept to be below the threshold in the short term and bet on the second period – hence the label 'hibernation'. Such a strategy is not viable in the 'game-over' and 'all in vain' scenarios – the game is over if $x_1 < 0$. This leads to extreme time discounting below the threshold: one needs to eat today in order for there to be a tomorrow.

3.4.3 The effect of future expectatoins

In all scenarios except 'all in vain', the effect of future expectations d on time preference is U-shaped (Fig. 3.4). The increasing discounting for large positive d is a known result: when tomorrow looks bright, one should prioritize present needs (Epper, 2015; Mell et al., 2021). By contrast, the increasing discounting for negative d is novel. It arises from the desperation threshold: if the threshold can still be reached today, but you strongly expect to fall below tomorrow, the best strategy is to enjoy the present. The effect is more pronounced in the 'hibernation' and the 'game-over' scenario: in these two cases, if below the threshold in the second period, it does not matter how much below. In this case, the agent tries to enjoy today, as tomorrow will be equally awful whatever she does. Of course, this strategy is impossible in the 'all in vain' scenario: if below the threshold in the second period, then the first period is also ruined, hence the monotonic effect of d in this setting (Fig. 3.4).

The U-shape of Fig. 3.4 is valid for any x_0 (see heatmaps in the SI), but the point of maximum patience shifts to $d \approx -x_0$, that is, when $\mathbb{E}(x_2) = 0$. In other words, patience peaks when the agent expects her basic needs to be in the balance in the long term, but not the short term. This result can be approached analytically in the 'additive penalty' and the 'hibernation' scenarios: for small $|x_0|$ and high w – that is, if the threshold is the main contribution on utility – we have the following approximation: $log(MRIS) \propto \frac{ln(2)}{2} + (d + x_0)^2 - 2x_0^2$, which is minimized when $d = -x_0$. This approximation also explains the parabolic shape of log(MRIS) around 0: downward-opening with respect to x_0 (Fig. 3.2) and upward-opening with respect to d (Fig. 3.4).

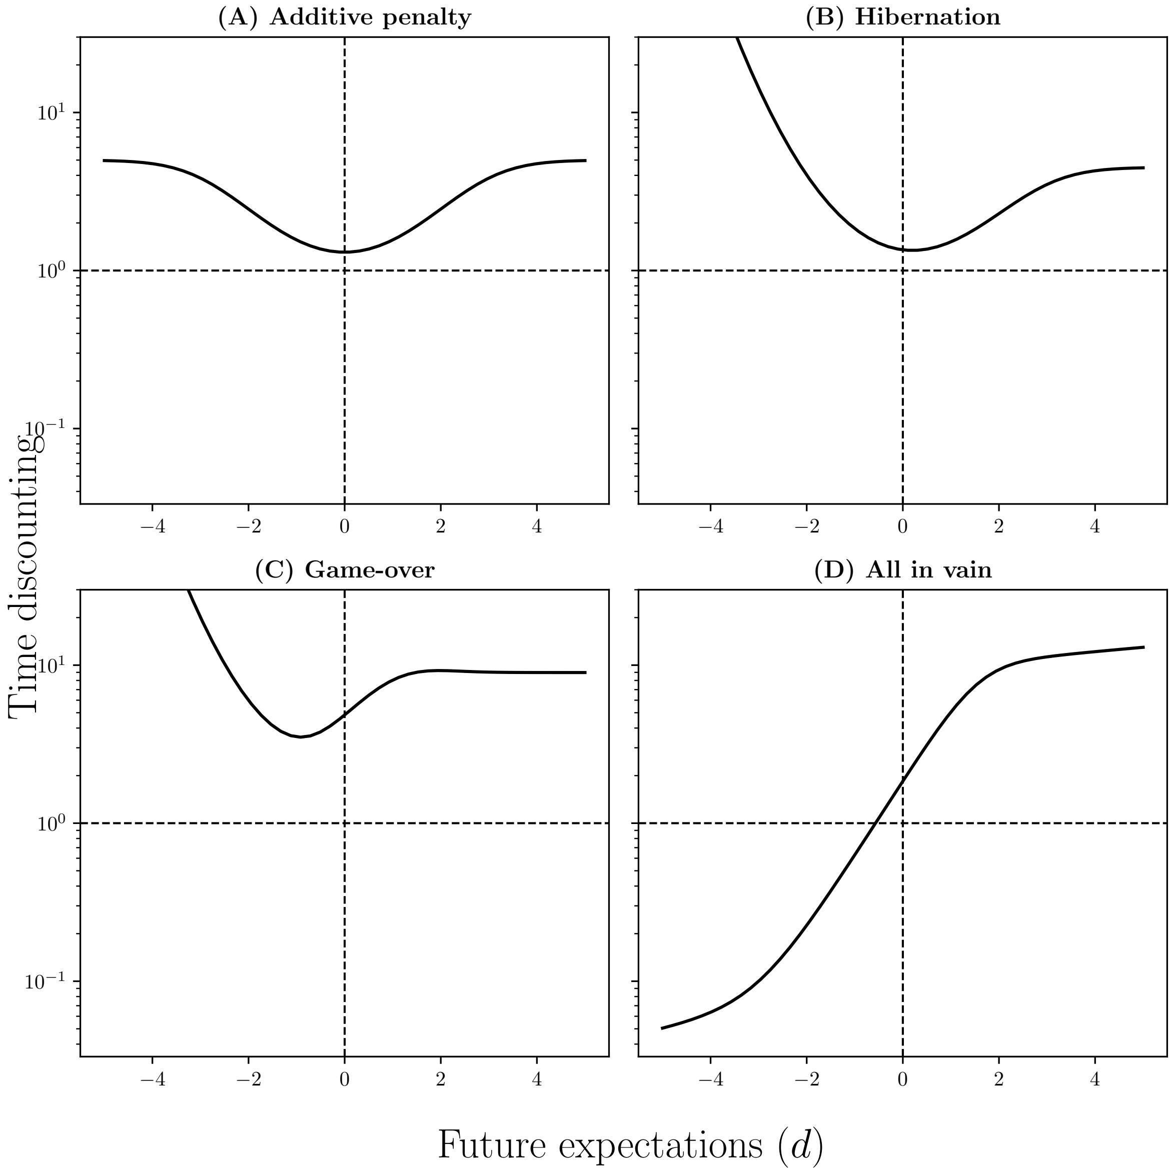

Figure 3.4: Time discounting depending on future expectations d, for $x_0 = 0$ and $\sigma = 1$. The vertical dashed line marks neutral future expectations, the horizontal one marks time neutrality. The y-axis is truncated at 1/30 and 30, for legibility.

3.5 Discussion

Time discounting in poverty is often justified on intuitive grounds (Epper, 2015; Fisher, 1930; Frankenhuis & Nettle, 2020; Haushofer et al., 2013; Mell et al., 2021; Sharma et al., 2023) by what we have called Earn's argument: in poverty, one 'needs to eat today', and thus prioritizes immediate needs. This paper began with a challenge to that idea: if one expects to also be in poverty tomorrow, then urgent needs today do not, by themselves, justify discounting the future. 'Need' is a folk psychological concept, and a clearer specification is needed. Based on the desperation threshold model, we propose that need corresponds to a steep region in the utility function – steeper than both above (where needs are met) and below (where they are not). We introduced four stylized scenarios of need, each producing distinct predictions. These divergences shows that Earn's argument is not trivial, and a formal model was needed. However, all four scenarios converge on one finding: near the desperation threshold, one should discount the future (Fig. 3.2). Counter-intuitively, this occurs not because the individual expects to need less resources in the future, but because she expects to be less likely to need them (Fig. 3.3). She definitely finds herself in need now, but is uncertain about whether she will find herself in need later; later, she may have either escaped poverty or fallen so far below that the additional resources then will be futile.

Two additional factors can amplify this effect. First, when agents hold directional expectations about their future – whether optimistic or pessimistic ($d \neq 0$), they expect to move further from the threshold in one direction or the other, and therefore discount the future more (Fig. 3.4). Second, when failing to meet the threshold carries irreversible consequences, as in the 'game over' and 'all in vain' scenarios, the agent focuses on short-term needs in order for there to be a tomorrow at all. This last result echoes previous models (Chavas, 2013), as well collection risk explanations, albeit in a state-dependent fashion. The two other effects (the effect of uncertainty and of $d \neq 0$) are however new, and direct consequences of the desperation threshold assumption.

When only the 'uncertainty' effect is present (Fig. 3.2A&B), our model predicts only a modest effect on time discounting: at best, $MRIS = \sqrt{2}$. This implies that one would trade 1€ in the first period for 1.41€ in the second. This result should not be interpreted literally: for technical and simplicity reasons (see Methods), our model compares two future time points rather than present versus future, with the second period being $\sqrt{2}$ times more uncertain than the first. In fact, maximum discounting is determined by the ratio of standard deviation of the resource levels at the two time periods. At the limit, if x_1 was known almost perfectly and we kept a discontinuous utility function, our model yields infinite discounting at the threshold in all four scenarios. This result also has an intriguing consequence if we vary the delay between time periods: in standard random walks or Brownian motion, variance increases linearly with time, which so the valuation of the reward is proportional to the square root of the delay. This is reminiscent of hyperbolic discounting, but with an even more extreme 'present bias': compared to exponential discounting, valuation drops heavily for short delays, and then very slowly. Thus, our model provides a possible rational explanation for non-exponential discounting and 'present-bias' which is commonly observed empirically, but often interpreted as 'time inconsistency' (Frederick et al., 2002).

Importantly, this effect occurs on both sides of the desperation threshold (Fig. 3.2). This contrasts with the implications of the desperation threshold on risk taking (de Courson et al., 2023; de Courson & Nettle, 2021), that predict a shift from extreme risk taking to extreme risk avoidance at the threshold. In other words, the DTM predicts that while poverty has a polarizing effect on risk taking, it uniformly increases time discounting. This aligns with empirical findings: studies find that poverty can increase or decrease risk taking, but the effect on time discounting is consistently positive (Haushofer & Salicath, 2023).

What real-world behaviors might our model explain? We argue that it provides a new possible rational account of decisions that alleviate need in the short-term, even if they lead to possibly greater need later. It applies particularly well to behaviors like acquisitive crime in contexts of pressing need (Jacobs & Wright, 1999; Rossmo & Summers, 2022; Topalli et al., 2015), or to payday loans uptake (Bertrand & Morse, 2011; Dobbie &

Skiba, 2013). It does not, as presented here, account for, for example, high consumption of 'temptation goods' in deprived populations (A. Banerjee & Mullainathan, 2010). To extend the model to cover those goods, we would have to posit needs that those goods serve, such as reputational or signalling needs. Alternatively, the consumption of such goods may be a maladaptive byproduct of increased discounting whose proper domain is need goods.

A second common prediction across our four scenarios is that when above the threshold but still vulnerable in the long run, one should be particularly patient, out of caution (Fig. 3.2). This intuition echoes the concept of 'prudence' and 'precautionary savings' in microeconomics. However, standard utility theory predicts precautionary saving across the social ladder, and in particular among individuals with low resources. In contrast, our model predicts that patience should peak among the 'middle class', that Ravaillon (Ravallion & Ravallion, 2016) defines as those "who are not considered 'poor' in a specific society but face a non-negligible chance of falling into poverty" (p.270). We are not aware of cross-sectional data supporting this prediction that the middle class is more patient than the upper class: on the contrary, Epper et al. (2020) finds that the negative association between resources and time discounting "exists through the wealth distribution" (p. 1177). However, our prediction echoes multiple theories in social sciences and economic history, which claim that patience and saving belong to so-called 'middle class values' (Doepke & Zilibotti, 2008; McCloskey, 2007; Veblen, 1899; Weber, 1905). This view is famously expressed by Weber (Weber, 1905), who argued that the 'protestant ethic' — emphasizing thrift and self-discipline — emerged among the upper-middle class, not the aristocracy. Doepke & Zilibotti (2008), studying 18th-century Britain, draws a more direct link to discounting behavior: "the middle class developed a system of preferences and values centered around parsimony, the work ethic, and delay of gratification," whereas "members of the upper class displayed a low propensity to save and accumulated debt, which suggests low patience" (p. 749–750). Our model also resonates with Sharma's recent findings (Sharma et al., 2023), which find a "polarizing effect of scarcity" (p. 1040) in both surveys and experiments. Specifically: "Perceptions of scarcity that threatened needs with shorter time horizons predicted a significantly increased preference for smaller, sooner outcomes [...], while perceptions of scarcity that threatened needs with longer time horizons predicted a significantly increased preference for larger, later outcomes" (p. 1040). We hope our model encourages further empirical work to rigorously test this prediction in contemporary populations.

However, the four scenarios diverge sharply in their predictions far below the threshold (Fig. 3.2). If falling below the threshold has irreversible consequences (as in the 'game over' and 'all in vain' scenarios), the model predicts extreme impatience. In the 'hibernation' case, the agent becomes maximally patient. In the 'additive penalty' scenario, she remains time-neutral. We do not argue that one of the scenario best captures empirical reality. Rather, each scenario may correspond to different real-world situations, depending on which basic need is threatened. For example, threats to survival such as starvation may be best modeled as an irreversible threshold, whereas more symbolic or social needs — such as Adam Smith's 'linen shirt' — may align with the logic of the 'hibernation' scenario, and rental eviction may fall somewhere in between. In any case, the specific predictions below the threshold have limited implications for real-world population patterns. Such states are likely rare and transient, as it one would seek external help or engage in risk-taking in this situation (de Courson, Frankenhuis, Gelder, et al., 2025). If these cases are indeed rare, all four scenarios converge on predicting a L-shaped association between resources and time discounting — high discounting in poverty, patience in the middle class, and time neutrality among the wealthy.

Our model also sheds light on how expectations about the future shape time preferences. In prior models of time discounting in poverty (Epper, 2015), and more generally in standard economic theory (Friedman, 1957), the effect of expectations is monotonic: the brighter one expects his future material situation to be, the more one prefers immediate resources. Yet in reality, extremely impatient behaviors such as acquisitive crime are often concentrated among populations with persistently pessimistic expectations (Anderson, 2000). In multiple large scale surveys, Gladstone & Pomerance (2025) actually finds an interaction effect: individuals with both low income and pessimistic beliefs save less than others (Fig. 3.4). Our model predicts a U-shaped effect of expectations. It can thus reconcile prior models — we also find that optimism can generate time discounting — with these empirical findings.

These findings have implications for public policy. Impatient behavior in poverty is frequently viewed as irrational and harmful both to individuals themselves and to society as a whole. Tt is often addressed as a unfortunate cognitive bias or failing, through interventions that aim to alter perceptions of the future (Rung & Madden, 2018). Our model suggests, instead, that high discounting is an appropriate response to material need, and therefore that directly addressing material need is the proper way of addressing discounting, while interventions on future orientation may be either ineffective or ethically questionable. Ensuring that everyone's basic needs are met would reduce impulsive behaviors and may yield positive externalities for society. It also makes a counterintuitive prediction: among middle-class individuals, anxiety about meeting future basic needs can lead to excessive saving — and, as a result, under-consumption or under-investment. An ambitious welfare state could therefore enable people with both adequate and inadequate incomes to better align consumption with their actual preferences, rather than having to adapt their behavior to the looming threat of falling below a desperation threshold.

3.6 Appendix

3.6.1 Mathematical derivations

Table 3.1: Parameters used in the model

Symbol	Interpretation	Type
x_0	Initial resource level	Parameter
x_1	Resource level in t_1	Random variable
x_2	Resource level in t_2	Random variable
σ	Resource noise between time periods	Parameter
d	Future resource expectations	Parameter
w	Utility jump at the threshold	Parameter
$U(.,.)$	Utility function	Function
$\mathbb{E}(U)$	Expected utility	Function
$\frac{d\mathbb{E}(U)}{d.}$	Marginal utility	Function
$MRIS$	Marginal rate of intertemporal substition	Outcome

3.6.1.1 Dynamics

The agent starts with a resource level x_0, and lives through two time periods. Between each time period, he is affected by a random perturbation, normally distributed with mean 0 and variance σ^2, with σ capturing the future uncertainty. In x_2, we also add a 'drift' parameter d, representing the agent expectation – if $d > 0$, she expects her situation to improve, while if $d < 0$, she expects it to degrade. We also add the possibility to add a small resource Δx_i in either of the time steps, to study the marginal effect of having more resources at some point in time. Note that these terms are only present in one time period: the term Δx_1 is absent in the second time period – we justify this choice in the Methods section of the main manuscript.

In other words,

$$x_1 = x_0 + \Delta x_1 + \epsilon_1$$

And,

$$x_2 = x_1 + \Delta x_2 + d + \epsilon_2 = x_0 + \Delta x_2 + d + \epsilon_1 + \epsilon_2,$$

with $\forall i \in \{1, 2\}, \epsilon_i \sim \mathcal{N}(0, \sigma^2)$

3.6.1.2 Additive penalty

In this scenario, the utility is the final state, minus a penalty for each period spent below the threshold 0.

$$U(x_1, x_2) = x_1 + x_2 - w * (x_1 < 0 + x_2 < 0)$$

Clearly, $x_1 \sim \mathcal{N}(x_0, \sigma^2)$ $x_2 \sim \mathcal{N}(x_0, 2\sigma^2)$

We have (denoting ϕ_σ the normal probability density of mean 0 and variance σ^2:

$$\begin{aligned}
\mathbb{E}(U) &= 2x_0 - w(p(x_1 < 0) + p(x_2 < 0)) \\
&= 2x_0 - w(p(\epsilon_1 < -x_0) + p(\epsilon_1 + \epsilon_2 < -x_0 - d)) \\
&= 2x_0 - w\left(\int_{-\infty}^{-x_0} \phi_\sigma(x)dx + \int_{-\infty}^{-x_0-d} \phi_{\sqrt{2}\sigma}(x)dx\right) \\
&= 2x_0 - w\left(\int_{-\infty}^{-x_0} \frac{1}{\sigma}\phi(\frac{x}{\sigma})dx + \int_{-\infty}^{-x_0-d} \frac{1}{\sqrt{2}\sigma}\phi\left(\frac{x}{\sqrt{2}\sigma}\right)dx\right)
\end{aligned}$$

To compute time discounting through the marginal rate of intertemporal substitution, we modify this formula by adding the option to add a small resource Δx_i in one of time steps. The formula becomes:

$$\mathbb{E}(U) = 2x_0 + \Delta x_1 + \Delta x_2 - w\left(\int_{-\infty}^{-x_0-\Delta x_1} \frac{1}{\sigma}\phi(\frac{x}{\sigma})dx + \int_{-\infty}^{-x_0-d-\Delta x_2} \frac{1}{\sqrt{2}\sigma}\phi\left(\frac{x}{\sqrt{2}\sigma}\right)dx\right)$$

Now, we compute the MRIS, which is defined as:

$$MRIS = \frac{\frac{dE(U)}{dx_1}}{\frac{dE(U)}{dx_2}}$$

By dx_i, we mean that we look at the marginal effect of the resource Δx_1 on utility.

In our case, differentiating under the integral and using the symmetry of the normal density, we get:

$$MRIS(x_0) = \frac{1 + \frac{w}{\sigma}\phi(\frac{x_0}{\sigma})}{1 + \frac{w}{\sqrt{2}\sigma}\phi(\frac{x_0+d}{\sqrt{2}\sigma})}$$

We can make a few observations: if $d = 0$, the MRIS is an even function ($\forall x_0 \in \mathbb{R}, MRIS(-x_0) = MRIS(x_0)$), symmetrical around the desperation threshold 0. It reaches a maximum at 0, that is, at the desperation threshold ($MRIS'(0) = 0, MRIS''(0) < 0$). This maximum is:

$$MRIS(0) = \frac{1 + \frac{w}{\sqrt{2\pi}\sigma}}{1 + \frac{w}{2\sqrt{\pi}\sigma}}$$

This maximum discounting value clearly increases with w (more generally one can prove that for all x_0 values $MRIS(x_0)$ increases with w). For very high values of w (a situation where utility approximates a step function around the threshold), we have:

$$MRIS(0) \to \sqrt{2}$$

For $d = 0$, we can derive analytically the condition for the agent to discount the future:

$$\begin{aligned}
MRIS > 1 &\iff \frac{dE(U)}{dx_1} > \frac{dE(U)}{dx_2} \\
&\iff 1 + \frac{w}{\sigma}\phi(\frac{x_0}{\sigma}) > 1 + \frac{w}{\sqrt{2}\sigma}\phi(\frac{x_0}{\sqrt{2}\sigma}) \\
&\iff e^{-x^2/2\sigma^2} > \frac{1}{\sqrt{2}}e^{-x^2/4\sigma^2} \\
&\iff e^{-x^2/4\sigma^2} > \frac{1}{\sqrt{2}} \\
&\iff \frac{-x^2}{4\sigma^2} > ln(\frac{1}{\sqrt{2}}) = \frac{-ln(2)}{2} \\
&\iff |x| < \sigma\sqrt{2ln(2)}
\end{aligned}$$

For extreme values of x_0, $\phi = o(1)$ whatever the variance, so $MRIS \to 1$ for $x \to \pm\infty$.

3.6.1.3 Hibernation

Now, utility is:

$$U(x_1, x_2) = (x_1 + w)(x_1 > 0) + (x_2 + w)(x_2 > 0)$$

Therefore,

$$\mathbb{E}(U) = \int_{-x_0-\Delta x_1}^{+\infty} (x_0+\Delta x_1+x+w)\phi(x)dx + \int_{-\infty}^{+\infty}\int_{-x_0-d-y-\Delta x_2}^{+\infty} (x_0+\Delta x_2+x+y+d+w)\phi(x)\phi(y)dxdy$$

Marginal utilities are computed in a similar way as in the previous scenario, using Leibniz rule:

$$\frac{d\mathbb{E}(U)}{dx_1} = \int_{-x_0}^{+\infty} \phi_\sigma(x)dx + w\phi_\sigma(-x_0) = \Phi_\sigma(x_0) + w\phi_\sigma(x_0) = \Phi(x_0/\sigma) + \frac{w}{\sigma}\phi(x_0/\sigma)$$

$$\frac{d\mathbb{E}(U)}{dx_2} = \int_{-x_0-d}^{+\infty} \frac{1}{\sigma\sqrt{2}}\phi(\frac{x}{\sigma\sqrt{2}})dxdy + \frac{w}{\sigma\sqrt{2}}\phi(\frac{x_0+d}{\sigma\sqrt{2}}) = \Phi(\frac{x_0+d}{\sigma\sqrt{2}}) + \frac{w}{\sigma\sqrt{2}}\phi(\frac{x_0+d}{\sigma\sqrt{2}})$$

Like in the first scenario, marginal utility in the second time period has the same shape as in the first, with only an increased variance. The formula has an intuitive interpretation. It is the probability to be above the threshold by that time (in which case you derive utility from extra resources), plus the density at the threshold, weighted by w (the utility at the threshold).

It is not possible to derive analytically a necessary and sufficient condition for $MRIS > 1$, but the one found in the first scenario can be shown to be a sufficient condition for positive x_0. Numerical results show that this actual region is slightly larger (see Fig. 3.2B in the main manuscript).

Indeed, if $x_0 > 0$, $\Phi_\sigma(x_0) > \Phi_{\sqrt{2}\sigma}(x_0)$, and as in the previous scenario, $\phi_\sigma(x_0) > \phi_{\sqrt{2}\sigma}(x_0) \iff |x_0| < \sigma\sqrt{2ln(2)}$.

Therefore, $0 < x_0 < \sigma\sqrt{2ln(2)} \implies \frac{d\mathbb{E}(U)}{dx_1} > \frac{d\mathbb{E}(U)}{dx_2} \implies MRIS > 1$.

For $x_0 \to +\infty$, in both cases $\Phi \to 1$ and $\phi \to 0$ (intuitively, one becomes sure to be above the threshold, and the density at the threshold vanishes). Clearly, $\frac{d\mathbb{E}(U)}{dx_1} \to 1$ and $\frac{d\mathbb{E}(U)}{dx_2} \to 1$, so $MRIS \to 1$.

For $x_0 \to -\infty$, in both cases Mills ratio implies that $\Phi = o(\phi)$: it becomes very unlikely to be above the threshold, so the dominant contribution to marginal utility comes form the possibility that the resources helps overcome the threshold. In this case, however, $\phi_\sigma = o(\phi_{\sqrt{2}\sigma})$: the normal distribution with a higher variance decays much slower. This implies that $\frac{d\mathbb{E}(U)}{dx_1} = o(\frac{d\mathbb{E}(U)}{dx_2})$, so $MRIS \to 0$.

For $x_0 \to +\infty$, clearly $P(x_1 > 0 \cap x_2 > 0) \to 1$ and the other terms become negligible, so $MRIS \to 1$.

3.6.1.4 All in vain

Now,

$$U(x_1, x_2) = (w + x_2)(x_1 > 0 \cap x_2 > 0)$$

Therefore,

$$\mathbb{E}(U) = \int_{-x_0-\Delta x_1}^{+\infty} \int_{-x_0-\Delta x_2-d-y}^{+\infty} (w + x_0 + \Delta x_1 + \Delta x_2 + x + y)\phi(x)\phi(y)dxdy$$

$$\begin{aligned}\frac{d\mathbb{E}(U)}{dx_1} &= 1 \cdot \int_{-x_0}^{+\infty} \int_{-x_0-d-y}^{+\infty} \phi(x)\phi(y)dxdy + \phi(x_0)\int_{-d}^{+\infty} (w + d + x)\phi(x)dx \\ &= P(x_1 > 0 \cap x_2 > 0) + \phi_\sigma(x_0)((w + d)\Phi_\sigma(d) + \sigma^2\phi_\sigma(x_0 + d))\end{aligned}$$

This formula is cumbersome, but can be interpreted. The first term is the probability to be able to derive utility from the resources. The second term is the density to be at the threshold in the first time period, weighted by several parameters. First, $w + d$, as $x_1 = 0$, $\mathbb{E}(x_2) = d$, so one expects to get a utility $w + d$. Second, $\Phi_\sigma(d)$, which is the probability to stay above the threshold in the second time period, given that $x_1 = 0$ (if one falls back below the threshold in the second period, having reached it in the first is useless). Last, one adds $\sigma^2\phi_\sigma(x_0 + d))$, which is the partial expectation of x_2 conditional on x_2 > 0: there is a chance that one gets more resources in the second period, which brings utility that one would not have had if not above the threshold in the first.

The marginal utility in the second period is:

$$\frac{d\mathbb{E}(U)}{dx_2} = 1 \cdot \int_{-x_0}^{+\infty} \int_{-x_0-y}^{+\infty} \phi(x)\phi(y)dxdy + w \cdot \int_{-x_0}^{+\infty} \phi(y)\phi(-x_0 - d - y)dy$$

The second member simplifies (after a change of variable):

$$\begin{aligned}\frac{d\mathbb{E}(U)}{dx_2} &= P(x_1 > 0 \cap x_2 > 0) + \frac{e^{\frac{-(x_0+d)^2}{4}}}{2\pi\sigma^2} \int_{-x_0}^{+\infty} e^{(\sqrt{2}y+(x_0+d)/\sqrt{2})^2/2} dy \\ &= P(x_1 > 0 \cap x_2 > 0) + w\phi_{\sqrt{2}\sigma}(x_0 + d)\, \Phi_{\sqrt{2}\sigma}(x_0 - d)\end{aligned}$$

This formula is simpler to interpret: the first term is still the probability to enjoy the extra resource because above the threshold in both periods, and the second is the density of x_2 in 0 (intuitively, the chance that the extra resource pushes one above the threshold), weighted by w (the utility at the threshold), and $\Phi_{\sqrt{2}\sigma}(x_0 - d)$ gives, intuitively, the probability to have been above the threshold in the first period given that one is right at the threshold in the second.

For very low values of x_0, we can look for equivalents of the terms. First, we get the equivalent for the probability to be above in both time steps. The intuition here is that the normal density decays so fast that the majority of the probability mass allowing survival is concentrated just above 0. So, if $x_1 > 0$ in spite of $x_0 << 0$, most likely $x_1 \sim 0$, and then we need $\epsilon_2 > -d$ to stay above the threshold in t_2. Formally, we split the integral and show that the rest becomes negligible. Here goes:

$$\begin{aligned}\mathbb{P}(x_1 > 0, x_2 > 0) &= \int_{-x_0}^{+\infty} \int_{-x_0-y}^{+\infty} \phi(x)\phi(y)\, dx\, dy \\ &= \int_{-x_0}^{+\infty} \int_{0}^{+\infty} \phi(x)\phi(y)\, dx\, dy + \int_{-x_0}^{+\infty} \int_{-x_0-y-d}^{-d} \phi(x)\phi(y)\, dx\, dy \\ &= \Phi(x_0)\Phi(d) + R(x_0)\end{aligned}$$

Now, inside the rest integral, we remark that $\forall x \in [-x_0 - y - d, -d], \phi(x) \leq \phi(-d) = \phi(d)$. We minimize the rest, and then we can make an explicit integration of the lower bound.

$$\begin{aligned}0 \leq R(x_0) &\leq \phi(d) \int_{-x_0}^{+\infty} \phi(y)(x_0 + y)dy \\ &= \phi(d)\left[x_0\Phi(x_0) + \phi(x_0)\right]\end{aligned}$$

Now, we use the Mills ratio: for very negative x, we have:

$$x\Phi(x) \sim -\phi(x)\left(1 - \frac{1}{x^2} + o(\frac{1}{x^2})\right)$$

Therefore,

$$R(x_0) \sim \frac{\phi(x_0)}{x^2} = o(\Phi(x_0))$$

Since we have:

$\Phi(x_0)\Phi(d) \leq \mathbb{P}(x_1\, 0, x_2\, 0) \leq \Phi(x_0)\Phi(d) + R(x_0)$

And $R(x_0) = o(\Phi(x_0)\Phi(d))$, the upper bound is equivalent to the lower bound, and by a sandwich theorem, we have:

$$\mathbb{P}(x_1 > 0, x_2 > 0) \sim \Phi_\sigma(x_0)\Phi_\sigma(d) \sim -\sigma^2 \Phi_\sigma(d)\frac{\phi_\sigma(x_0)}{x_0}$$

For $w > 0$, the other member of $\frac{d\mathbb{E}(U)}{dx_1}$ is clearly dominant, and we have:

$$\frac{d\mathbb{E}(U)}{dx_1} \sim w\frac{\phi(x_0/\sigma)}{\sigma}\Phi(d/\sigma)$$

As for $\frac{d\mathbb{E}(U)}{dx_2}$, we can compute an equivalent of the second member, using Mills ratio and simplifying:

$$\frac{w}{\sqrt{2}}\phi_{\sqrt{2}\sigma}(x_0 + d)\Phi_{\sqrt{2}\sigma}(x_0 - d) \sim -\frac{w}{\sqrt{2}}\phi(x_0/\sigma)\frac{\phi(d/\sigma)}{x_0}$$

Let us assume that w is large enough, so that this term dominates the other term – otherwise the terms are of the same order and we can combine them, leading to a similar behavior. In this case, we obtain:

$$\begin{aligned} MRIS = \frac{\frac{d\mathbb{E}(U)}{dx_1}}{\frac{d\mathbb{E}(U)}{dx_2}} &\sim \frac{w\frac{\phi(x_0/\sigma)}{\sigma}\Phi(d/\sigma)}{-\frac{w}{\sqrt{2}}\phi(x_0/\sigma)\frac{\phi(d/\sigma)}{x_0}} \\ &\sim -x_0\frac{\Phi(d/\sigma)}{\sqrt{2}\sigma\phi(d/\sigma)} \end{aligned}$$

For $x_0 \to +\infty$, clearly $P(x_1 > 0 \cap x_2 > 0) \to 1$ and the other terms become negligible, so $MRIS \to 1$.

3.6.1.5 Game-over

In this scenario, the utility function is:

$$U(x_1, x_2) = (x_1 + w)(x_1 > 0) + (x_2 + w)(x_1 > 0 \cap x_2 > 0)$$

Therefore,

$$\mathbb{E}(U) = \int_{-x_0-\Delta x_1}^{+\infty}(x_0{+}\Delta x_1{+}x{+}w)\phi(x)dx + \int_{-x_0}^{+\infty}\int_{-x_0-\Delta x_2-y}^{+\infty}(w{+}x_0{+}\Delta x_2{+}x{+}y{+}d)\phi(x)\phi(y)dxdy$$

Marginal utility in the first period is:

$$\frac{d\mathbb{E}(U)}{dx_1} = \Phi_\sigma(x_0) + w\phi_\sigma(x_0) + \phi_\sigma(x_0)((w + d)\Phi_\sigma(d) + \sigma^2\phi_\sigma(d))$$

The computation of marginal utility in t_2 is the same as the previous model:

$$\frac{d\mathbb{E}(U)}{dx_2} = P(x_1\ 0 \cap x_2\ 0) + w\phi_{\sqrt{2}\sigma}\,(x_0 + d)\,\Phi_{\sqrt{2}\sigma}(x_0 - d)$$

Using the same reasooning as in the previous scenario, we can show that the MRIS becomes approximately linear for $x_0 \to -\infty$, but with a steeper asymptotic slope. For high w, we have:

$$MRIS \sim -x_0 \frac{1 + \Phi(d/\sigma)}{\sqrt{2}\sigma\phi(d/\sigma)}$$

We note that for $d = 0$, the slope is precisely three time steeper than in the 'all in vain' scenario. For $d >> 0$, it is two times steeper. For $d << 0$, it is infinitely steeper.

3.6.2 Sensitivity analysis

The main text shows the effect of x_0 and d on time discounting. Here, we present the effect of the other parameters (σ and w).

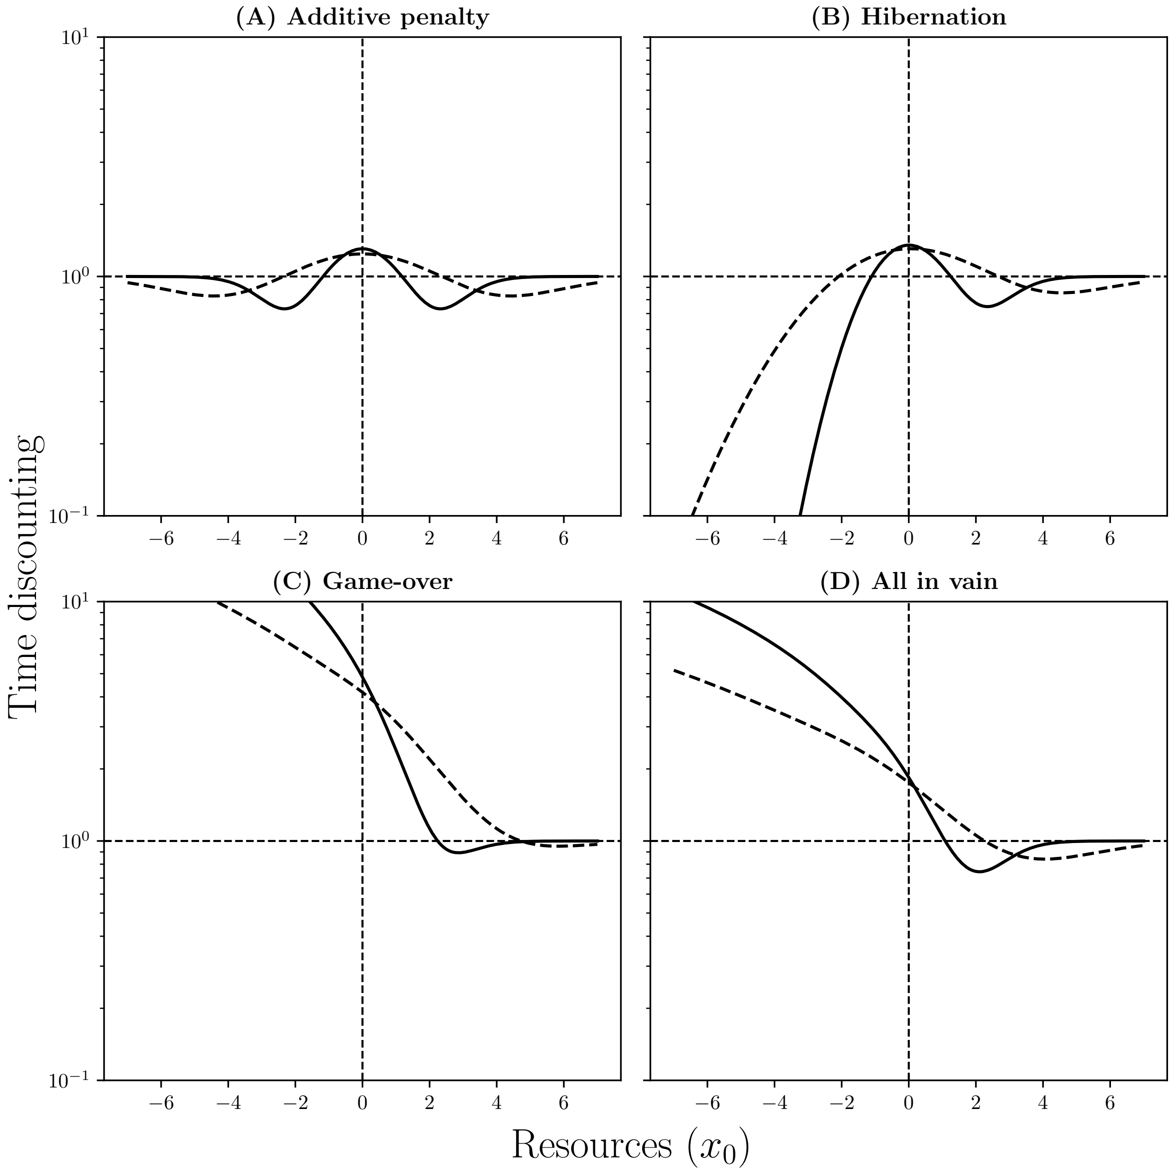

Figure 3.5: Time discounting depending on resources, for high and low resource noise between periods. The solid line denotes $\sigma = 1$, the dashed line denotes $\sigma = 2$. The y-axis is truncated at .1 and 10 for legibility.

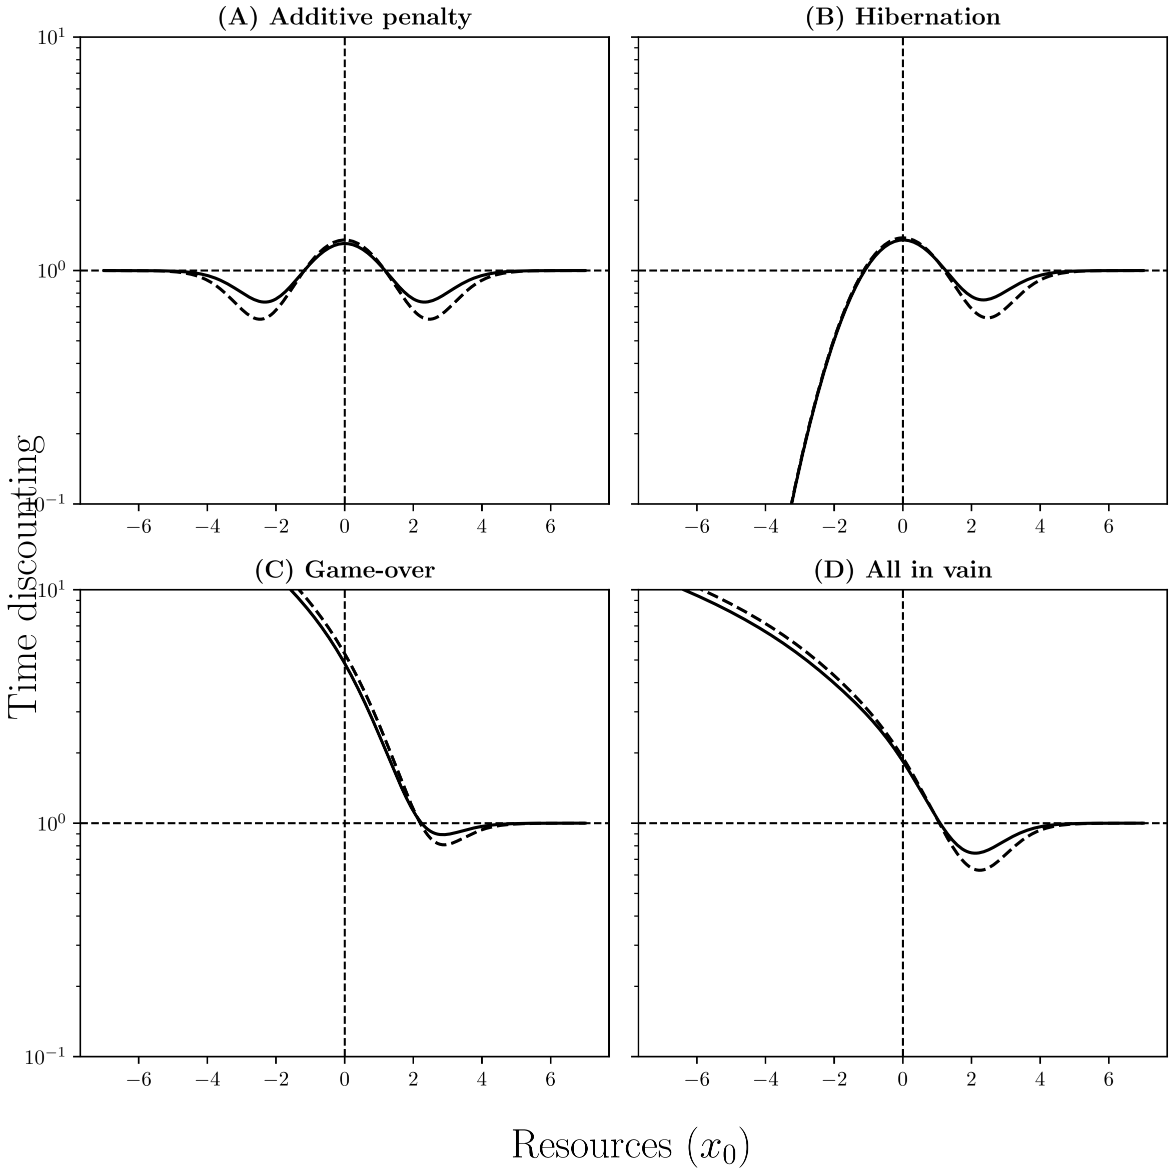

Figure 3.6: Time discounting depending on resources, for high and low utility jump at the desperation threshold. The solid line denotes low jump ($w = 10$), the dashed line denotes high jump ($w = 20$). The y-axis is truncated at .1 and 10 for legibility.