

**Uncertain survival and time discounting: intertemporal consumption plans
for family trusts.**

Stephen E. Satchell

Trinity College, University of Cambridge

Trinity Street

Cambridge CB2 1TQ United Kingdom

Susan Thorp*

School of Finance and Economics, University of Technology, Sydney

PO Box 123

Broadway NSW 2007 Australia

Email: susan.thorp@uts.edu.au

Abstract:

We derive expressions for optimal consumption for family trusts with random wealth and uncertain survival. Using UK birth statistics and the theory of branching processes, we compute size and survival probabilities for single and multiple-branch families. Survival for a single-branch family is approximated by a Pareto distribution and consumption policies exhibit decreasing discount rates but multiple-branch families use non-monotonic discount rates. When trust distributions depend on the number of beneficiaries rather than the survival of the whole family unit, spending paths depend on expected membership and the elasticity of intertemporal substitution. We report examples of consumption paths for a range of family trusts with CRRA preferences.

JEL Classification: G0, D91, D81

Keywords: intertemporal choice, family extinction, hyperbolic discounting

⁰*Corresponding author

1 Introduction

Many problems in economics and finance require planning over long time horizons and in all such problems the discount rate is critical. Researchers from Ramsey (1928) to Stern (2006) have recognized that choosing a discount rate is not always a disinterested decision but frequently represents some amalgamation of economic ‘science’ with inter-generational ethics or politics.

While it is analytically convenient for economists to choose non-stochastic boundaries (either fixed or infinite) for multi-period problems, more often than not a problem’s inherent uncertainty extends beyond, say, random investment payoffs or uncertain income, to the horizon itself. If the planning horizon is stochastic, discounting cannot arbitrarily be fixed at some unobserved level of impatience: it must be treated as a function of the probability density of the horizon. Further, stochastic horizon problems are ubiquitous. The simplest individual consumption problems are subject to uncertainty over the length of life, and the same is true of the majority of plans for firms, financial institutions, governments and societies.

Time-varying discounting has sometimes been used to harmonize observed patterns of behaviour with the predictions of theory and attributed to weak ethics, a lack of altruism on the part of the decision-maker or an intrinsic psychological preference for near-term gratification. For example, hyperbolic discounting is a well-known explanation for anomalous savings behaviour (see, for example, Phelps and Pollack, 1968; Laibson, 1998). However a more recent literature (Sozou 1998; Dasgupta and Maskin 2005; Bommier 2006) has shown that regardless of psychological or ethical considerations, horizon uncertainty alone may be a sufficient explanation for some examples of hyperbolic discounting.

Here we offer a new explanation for time-varying discounting that does not rely on arguments relating to the failure of altruism or fundamental tastes, requiring only rational uncertainty over the long-term survival of the planning entity. Since, as we show below, the empirical survival function of a multi-generation family from a single progenitor has a hazard rate that declines with the age of the family, hyperbolic discounting may apply to the planning problem of the family trust. In the case where the family begins with several branches, the discount function is non-monotone, with hazard rates that increase for several generations before declining hyperbolically. Further, if the family does not pass all wealth to a single heir but distributes according to the number of surviving members, then the discount function depends on expected family growth and the elasticity of intertemporal substitution.

Since we assume that trusts are risk neutral over survival and that subjective time preference is constant, discount rates are not revised with the passage of time (unless new information on survival arrives) and plans are time-consistent.

We use numerical optimization methods calibrated to UK birth statistics and a representative investment model to estimate optimal spending paths for a family trust begun by a single progenitor, by three progenitors, and for a trust which distributes according to the number of members, where all are maintained along one gender line. (We make use of the theory of branching processes to calculate the probability of family size and family extinction at each generation.) For a foundation or trust distributing without regard to family size and expecting real investment returns at 4.75% each year, the ideal annual spending rate for the single progenitor trust begins around 2.1% of wealth, compared with the infinite horizon optimal rate of 1.54%, and declines slowly as the generations pass. For a trust distributing according to membership with a power utility parameter of 2, the spending rate begins low, at just above 1%, before gradually increasing above the infinite horizon rate and then declining.

2 Literature

In his pioneering work on aggregate savings, Ramsey (1928) asserts that any positive discounting of the future is ‘ethically indefensible and arises merely from the weakness of the imagination’. Ramsey actually relents from his uncompromising view by using a non-zero discount rate in the analysis which follows this statement, but he does rule out the possibility of ‘savings being selfishly consumed by a subsequent generation’. Others are less optimistic about the strength of imagination than Ramsey, allowing that the current generation could be less-than-perfectly altruistic towards future generations. Phelps and Pollack (1968), for example, consider a multi-generation model in which consumption in period t is discounted by bv^t , where v is the rate of time preference and b ($0 < b, 1$) represents the current generation’s altruism. The closer b is to one, the more concerned is the current generation about the welfare of future generations. They recognize that if succeeding generations have these same quasi-hyperbolic preferences but cannot control the savings behaviour of their descendants, the outcome is a Nash-equilibrium where saving is lower than the Pareto-optimal level. The current generation rationally consume faster than the Pareto-optimal rate in an effort to limit over-spending by their children and grandchildren.

Similar outcomes can occur when a buffer-stock consumer plays an intra-personal

game with future ‘selves’ (Laibson, 1998; Harris and Laibson, 2001). The ‘current self’, a hyperbolic discounter, expects ‘future selves’ to over-consume relative to the current self’s preferences. The effective rate of impatience in this case depends on future scarcity, is stochastic, and endogenous to the model. Laibson (1998) and Laibson et al. (1998) argue that this type of discounting can explain savings behaviour that seems inconsistent with a standard exponentially discounted model, results in welfare losses, and can help explain the documented tendency among people and animals to discount near events more than distant events (Loewenstein and Thaler, 1989; Ainslie, 1974).

On the other hand, more recent explanations for decreasing rates of time preference arise from horizon uncertainty rather than preferences for short-term gratification. Sozou (1998) looks at an uncertain future payoff of fixed size where the probability of receiving the payoff in any future period is determined by an exponential survival function with constant hazard rate. However if the consumer does not know the true underlying value of the hazard rate, but holds a prior belief that it is exponentially distributed, then he or she will compute a hyperbolic discount function by Bayesian updating.

Dasgupta and Maskin (2005) argue that while both a declining hazard rate and Sozou’s analysis can produce hyperbolic discounting behaviour in the sense of decreasing impatience, these cannot explain preference reversals, where a consumer switches from one course of action to another simply because of the passage of time, or time-inconsistent behaviour. Dasgupta and Maskin’s own explanation for hyperbolic discounting rests on uncertainty over when, rather than just whether, a payoff will be realized. This additional dimension of uncertainty can result in a preference reversal, and if learning is also needed, can generate time-inconsistency.

Life-cycle consumption and investment problems frequently model horizon uncertainty as stochastic individual mortality with non-constant hazard rates. (See, for example, Yaari 1965; Hubbard Skinner and Zeldes 1995; and for a survey of recent actuarial literature, Pitacco 2004). In a new insight, Bommier (2006) introduces risk preferences over the length of life to an intertemporal consumption model, demonstrating that standard treatments of individual survival uncertainty assume risk neutrality over lifetimes. Among other general results, he shows that hyperbolic discounting will be exhibited by individuals who have ‘no pure time preference’ but have a hyperbolic risk aversion to length of life, so that attitudes to future mortality risk rather than future consumption risk create near-term impatience.

Here we extend this line of research by deriving the general result that rational agents facing uncertainty over a long-term planning horizon will exhibit hyperbolic discounting

when survival is Pareto distributed and agents have risk neutral preferences over survival. Pareto distributed survival implies declining hazard rates and consequently decreasing impatience over more distant events. Further, we demonstrate that family survival from a single progenitor may be well represented by a Pareto density. We also investigate related questions of family survival: when more than one branch of the family survives, hazard rates are not monotonically declining and discounting functions are not uniformly hyperbolic, and when distributions depend on the number of family members rather than just survival, discounting reflects fertility and the elasticity of intertemporal substitution.

Models where hyperbolic discounting arises from the psychological tastes of the decision-maker predict preference reversals and dynamic inconsistency (Laibson 1998, Bommier 2006). Inconsistencies arise because the discount rate that was applied to a specific future date changes to a higher rate as the date moves closer to the decision period. By contrast in our analysis, the discount factor depends on the absolute age of the family rather than its age relative to the current period. Consequently, in the absence of new information about future survival probabilities, the survival function of the family is not revised as time passes and plans will be time-consistent (assuming risk neutrality with respect to survival). However if the family receives new information on fertility status, either on the number of survivors or on their prospects for having children, the trust may revise estimates of the survival density and consequently change future consumption plans. This does not imply time-inconsistency in the usual sense, just path-dependency in the optimal plan.

Our discussion so far has treated the problem of maximizing family or multi-generational utility as equivalent to maximising an individual's utility over an uncertain lifetime. This approach is typical of unitary models of intergenerational transfers where the 'head', 'altruist' or 'dictator' allocates consumption among current and future members of the family so that the welfare of the family is indistinguishable from the welfare of the head (Becker 1974; 1981). For the analysis of a family trust, the unitary model provides an important special case. For 'head' we can substitute the label 'establisher', whose utility function is embodied in the trust deed and executed over the life of the trust by the trustee. Indeed, one reason for creating a trust is to ensure that the preferences of the original family head are put into effect, overriding any deviation motivated by the preferences and strategies of future beneficiaries. If co-operative bargaining could produce the same allocation arrangement (Manser and Brown 1980; McElroy and Horney 1981) or if it could be ensured by family social capital or a self-enforcing family constitution via exchange (Cigno 1993; 2006; 2007), then the trust deed would be redundant. The

fact that trust deeds exist is evidence that the family is a place of both ‘conflict and cooperation’ (Xu 2007).

Even so, empirical evidence shows that most people plan to divide their estates equally among their children (Light and McGarry 2004) and it is worth investigating the optimal plans of a trust that cares about more than a single heir. If total utility is comprised of the utility of current and future generations, in the style of Becker and Tomes (1986) or Becker and Barro (1988), and where consumption is divided equally among children in each time period and then aggregated, not only survival, but also the expected number of family members, matters to the optimal disbursement of trust funds. In this case, discounting depends on expected fertility as well as simple survival.

The theory of discounting under horizon uncertainty is outlined in section 3. In section 4 we estimate the survival function of a representative UK family beginning with an individual progenitor: we find that the Pareto distribution is a good fit to current fertility data for this case, that extinction is certain at observed birth rates, and that the mean survival of a UK family is about six generations. We also estimate the survival function of a family with more than one original branch (k progenitors) and derive expressions for the expected path of family size over time. Numerical estimates of optimal spending paths (section 5) for a family trust from a single progenitor using the estimated hyperbolic discounting function are shallow curves, always above the infinite-horizon spending rate, whereas for the multi-branch family they are non-monotonic. Further, when distributions depend on the number of surviving members, the pattern is reversed, with spending rates rising over time initially before decreasing later. Section 6 concludes.

3 Discounting under survival uncertainty

Our problem is to generalize the model of optimal drawdown for an infinitely-lived entity facing uncertain investment returns (Ingersoll 1987; Korn and Korn 2001) to include the case where the survival of the entity is also uncertain (Dasgupta and Maskin 2005; Bommier 2006). For a family trust, consumption stands for payments to current family beneficiaries, funded from an investment portfolio.

In most common law jurisdictions a family trust deed is invalid if it attempts to tie up wealth for the benefit of generations not yet in existence.¹ The common law ‘rule against perpetuities’ (Burke 1976), or codified law relating to the same issue, usually

¹We thank Mr Vincent Taubman of TD Asset Management for drawing our attention to some of the legal constraints on trust deeds and trustees.

requires the interest (assets) in the trust to be passed to beneficiaries within 80 or 90 years. However some US jurisdictions allow large dynastic trusts to escape the rule and exist for much longer (or perpetually) before vesting. In our analysis we assume either that the trust is exempt from the rule against perpetuities and can create a trust deed for future generations, or that the trust continues under a ‘rolling’ deed which the family voluntarily recreates at each generation. This latter assumption allows the family to review its fertility status and survival distribution at any generation and revise the trust deed accordingly.

We consider two cases: one where the family has a single utility function and another where utility is an aggregation of the utilities of individual beneficiaries. For the unitary utility case, we treat family extinction as a random time. When family utility is an aggregation of the utilities of individual beneficiaries, the number of surviving family members is intrinsic to the objective function, and since the trust terminates when membership goes to zero, we consider only family size as a measure of extinction. Although the aggregate utility function described below in equation (13) introduces one form of risk aversion over family survival (proxied by the number of family members), the more general treatment of risk aversion over length of life expressed in Bommier (2006) could also be applied to family survival. This is a much more complex problem which we leave for future work.

We do not specify the distribution of investment returns, except that they are assumed independent and identically distributed (i.i.d.). The trust is extinguished when the family ceases, so we treat residual trust funds as having no utility value when the family is not alive to enjoy them.

3.1 Unitary utility

We begin with the case where the distribution of the trust funds among the family matters only in aggregate, as would be the case where the trust deed requires inheritance to pass to one heir, or where the trust treats the consumption of individuals in the family at time t as perfectly substitutable. It follows that consumption of the beneficiaries in each period can be treated as a single sum and horizon uncertainty relates to the survival of the family as a whole.

Let T be the random time the family survives. We treat the survival time as a continuous random variable and denote as $pdf(t)$ the probability density of T , the extinction density with distribution function $F(t)$, and $\bar{F}(t)$ its complementary distribution func-

tion. It follows that $\frac{\partial \bar{F}(t)}{\partial t} = -pdf(t)$, and $\bar{F}(0) = 1$, $\bar{F}(\infty) = 0$, so that the family survives almost surely in period zero but eventual extinction is inevitable. We discuss the inevitability of extinction further below.

The trust aims to maximize expected utility for as long as the family survives, where utility is derived from consumption (the distributions of funds) out of stochastic wealth. The family has an individual utility function, an assumption that we will relax below. Let the utility of consumption be $\bar{U}[C(t)] = U[C(t)]h(t)$ where $h(t)$ is some positive discount function expressing general impatience; possibly $h(t) = 1$, and for now we assume $h(0) = 1$. If we write \mathbb{E}_C to mean expectation over consumption the value of expected utility conditioning on survival until time t is:

$$L(t) = \mathbb{E}_C \left(\int_0^t \bar{U}[C(s)] ds | T = t \right), \quad (1)$$

and

$$L = \int_0^\infty L(t) pdf(t) dt \quad (2)$$

is the unconditional value.

Now, integrating by parts,

$$\begin{aligned} L &= \int_0^\infty \mathbb{E}_C \left(\int_0^t \bar{U}[C(s)] ds \right) pdf(t) dt \\ &= - \left[\mathbb{E}_C \int_0^t \bar{U}[C(s)] ds \bar{F}(t) \right]_0^\infty + \mathbb{E}_C \left(\int_0^\infty \bar{F}(t) \bar{U}[C(t)] dt \right), \end{aligned} \quad (3)$$

using Leibnitz's rule, and simplifying, we arrive at

$$L = \mathbb{E}_C \left(\int_0^\infty \bar{F}(t) \bar{U}[C(t)] dt \right), \quad (4)$$

by using the fact that the first expression in (3) is zero when $\mathbb{E}_C \left(\int_0^0 \bar{U}[C(s)] ds \right) = 0$. If survival is exponentially distributed with a constant hazard rate $\bar{F}(t) = \exp(-\lambda t)$, and if general impatience is constant so that $h(t) = \exp(-pt)$ where p is the impatience parameter, we recover Blanchard's (1985) result,

$$L = \mathbb{E}_C \left(\int_0^\infty \exp(-(p + \lambda)t) U[C(t)] dt \right). \quad (5)$$

In other words, uncertainty over family survival simply increases impatience by a constant

hazard rate, raising consumption permanently above the optimal rate for an infinitely-lived dynasty.

However this analysis can also deal with hyperbolic discounting and quasi-hyperbolic discounting. Define hyperbolic discounting as a discount function

$$D(t) = (1 + \beta t)^{-\gamma/\beta}; \quad \beta > 0, \gamma > 0, \quad (6)$$

and quasi-hyperbolic discounting by

$$D(t) = bv^t; \quad 0 < b < 1, 0 < v < 1. \quad (7)$$

If we arbitrarily set general impatience at zero so that $h(t) = 1$, then the survival function is

$$\bar{F}(t) = (1 + \beta t)^{-\gamma/\beta}, \quad (8)$$

and

$$pdf(t) = \gamma(1 + \beta t)^{-(\gamma/\beta+1)} \quad (9)$$

is the density of family extinction, whilst for quasi-hyperbolic discounting the survival function is

$$\bar{F}(t) = bv^t,$$

and the density of family extinction is

$$\begin{aligned} pdf(t) &= -\ln(v)bv^t \\ &= \ln\left[\left(\frac{1}{v}\right)^b\right]v^t. \end{aligned} \quad (10)$$

The first density (9) can be thought of as Pareto, where the family survives almost surely in the initial period $\bar{F}(0) = 1$, but is extinct in the limit $\bar{F}(\infty) = 0$. The second (10) is not normalized in that $\bar{F}(0) = b$, and $\bar{F}(\infty) = 0$. If $v = \exp(-\lambda)$, $\lambda > 0$, and $\bar{F}(t) = b\exp(-\lambda t)$, this is similar to (5) and corresponds to the case where a proportion of families, $(1 - b)$, die initially at $t = 0$; so $\bar{F}(0) = b$. However, we could also have the case of a degenerate extinction probability in that $\bar{F}(t) = (1 - \rho) + \rho\bar{F}^*(t)$ where $\bar{F}^*(t)$ has the usual property ensuring extinction in the limit, $\bar{F}^*(\infty) = 0$. This brings about no important changes but allows us to incorporate $(1 - \rho)$, a finite probability that the family will last forever.

3.2 Individual utility

An alternative representation treats the family size as a non-negative process $N(t)$ with an absorbing state at $N(t) = 0$, and views utility as an aggregation of the individual utility of trust beneficiaries, $\tilde{U}\left(\frac{C(t)}{N(t)}\right)N(t)$. The probability of any future event $N(t)$ is known but not controlled by the trustees or the beneficiaries. In this specification family size determines extinction, that is, the family is extinct when the $N(t)$ process reaches zero for the first time and we simply consider the path of expected consumption and fertility as encompassing the relevant aspects of the family's survival, where

$$L = \mathbb{E}_{C,N} \left[\int_0^\infty \tilde{U}\left(\frac{C(s)}{N(s)}\right) N(s) ds \right]. \quad (11)$$

Particular choices of $\tilde{U}(\cdot)$ lead to particular solutions. We shall assume a power or constant relative risk aversion (CRRA) form, the same for each trust beneficiary, and constant across generations so that

$$\tilde{U}(C(t)) = \frac{C(t)^{1-\alpha}}{1-\alpha} \text{ for } \alpha \geq 0. \quad (12)$$

Hence

$$\tilde{U}\left(\frac{C(t)}{N(t)}\right) N(t) = \frac{C(t)^{1-\alpha}}{1-\alpha} N(t)^\alpha \quad (13)$$

and this specification leads to a 'Cobb-Douglas' version of utility with constant returns to scale. So the objective function can be written as

$$L = \left[\int_0^\infty \frac{\mathbb{E}_{C,N} [C(s)^{1-\alpha} N(s)^\alpha]}{1-\alpha} ds \right]. \quad (14)$$

In numerical examples below we model the distribution of the process $N(t)$ as known and outside the control of the trust beneficiaries. Becker and Barro (1988) analyse a similar problem under certainty over returns to wealth and where fertility and consumption are choice variables. Our trustee can choose only consumption but returns to wealth and family size are both (independent) random processes. And since, like Becker and Barro, we treat the path of $N(t)$ as predictable from $t = 0$ and make the preferences of the family (trustee) constant, the optimal path for distributions will be time-consistent if general impatience is also constant.

4 Estimating family extinction

We now apply this analysis by estimating the density function of survival for a representative UK family using current fertility data.

Theories of family extinction and the related literature on branching processes are associated with Sir Francis Galton, who posed the problem of the survival of aristocratic surnames in 1873 (see Harris, 1963; Kendall, 1966). Despite the fact that this is a standard problem in population studies, it appears that few empirical estimates of family extinction are available (Albertsen, 1995). But the question remains interesting in the light of changing fertility patterns over the past century where extinction probabilities have increased in many countries as fertility has declined. In the UK for example, the number of children of either sex being born to each woman (total fertility rate) was around 1.79 in 2005, which is below the long-term replacement rate of 2.1 required to maintain a stable population and much less than the 1960 peak total fertility rate of 2.95. Similarly, the OECD-30 average fertility rate in 2004 was 1.6, with birth rates of below 1.3 children per woman in Japan and some European countries, and as low as 1.2 in Korea (OECD 2007). Hence family survival prospects in OECD countries are now much lower than they were thirty years ago.

Early studies of family survival are rare, but those that do exist estimate limiting extinction probabilities less than one reflecting higher fertility rates in the past. For example, Lotka (1931) published an estimate of 0.8797 for the probability of male line extinction for the US white population of 1920, and Keyfitz (1968) calculated the likelihood of female line extinction at 0.8206 using 1960-61 US data, along with similar calculations for Hungary, Israel, Mexico and Japan. Hull (1998) reconsidered Lotka's calculations in the context of a population with two sexes and concluded that, under some restrictions over the availability of partners to the males of concern, the extinction probability lay in the range $(0.856, 0.992]$, not greatly different from Lotka's original estimate.

Our estimation follows the method of Keyfitz (1968): beginning with the official UK statistics on the distribution of women born in 1960 by number of live births at age 45, we adjust this probability to the number of daughters (assuming that male and female births are equally likely) and then compute the probability that the female line becomes extinct in the limit along with the probability that the line becomes extinct at or before any particular generation. We cannot be sure that the same distribution of birth probabilities applies to the male line, since paternity data collected in the UK are incomplete and there is no comparable table of birth probabilities for men, but it is plausible that a dynasty

which passes its wealth through sons rather than daughters might face similar survival probabilities. (Modern families may be less likely to be concerned with surname survival and more likely to pass wealth to sons and daughters than to favour either gender. This more general problem requires a different survival model with specific limits on partnering and population growth, and we leave it to future research.)

In addition, we assume that the group of families who create inherited trusts have fertility patterns the same as the 1960 birth cohort of women and that this pattern remains constant into the future. While there might be reasons to assume higher fertility among wealthier families because of better health prospects, lower fertility is also possible because of the well-documented tendency for more educated women to begin families later, resulting in fewer children (Rendall et al. 2005). In addition, fertility rates are not constant through time, having been declining in many countries since the middle of last century. For some countries including the UK, this decline seems to have slowed or reversed recently (see Office for National Statistics 2006, Table 1.4). For England and Wales, for example, the average number of live daughters a woman of child-bearing age of a particular cohort can expect to have in her lifetime (Gross Reproduction Rate) rose to 0.88 in 2005 from a low point of 0.81 in 2000. Since it is difficult to predict the long-term path of wealthy family fertility, we work with current population averages and project them into the future.

Further, we treat the fertility outcome as a random draw from the distribution of family size per female; we do not allow the mother to choose the number of her children. Survey analysis (OECD 2007) reports that women in most developed countries, including the UK, have fewer children than they want, on average, which is evidence that family size is not always a choice variable.

4.1 Family size and survival along a single branch

In 1930, two Danish mathematicians, Steffensen and Christensen, separately and simultaneously solved Galton's problem of surname survival. (Their work was reprinted in English in 1995, see Steffensen 1995; and Albertsen and Kristensen 1995. For a general introduction to branching processes see Taylor and Karlin 1998, chapter III.) Suppose that a mother (father) produces a random number ξ of daughters (sons) with probability distribution

$$\Pr \{\xi = k\} = a_k \quad \text{for } k = 0, 1, \dots \tag{15}$$

where $a_k \geq 0$ and $\sum_{k=0}^{\infty} a_k = 1$. In what follows we treat the a_k as known and constant through time. The family size process $\{G_g\}$ at generation g is a Markov chain where each G_g is a random sum of the offspring produced by the previous generation,

$$G_g = \xi_1^{(g-1)} + \xi_2^{(g-1)} + \dots + \xi_{G_{g-1}}^{(g-1)}. \quad (16)$$

where $\xi_j^{(g-1)}$ is the number of children born at generation g to offspring j , $j = 1, \dots, G_{g-1}$, from the previous generation, $g - 1$.

Following Taylor and Karlin (1998) we can write the probability generating function for ξ at the first generation as

$$\phi_1(s) = \mathbb{E}(s^\xi) = \sum_{k=0}^{\infty} a_k s^k, \text{ for } 0 \leq s \leq 1 \quad (17)$$

where the probability of k daughters is

$$a_k = \frac{1}{k!} \left. \frac{d^k \phi(s)}{ds^k} \right|_{s=0}.$$

The distribution of family size (measured by number of children of one gender) at any generation G_g can be derived from the probability generating function at generation g where $\phi_g(s)$ is the outcome of the iteration

$$\phi_g(s) = \phi_1[\phi_{g-1}(s)]. \quad (18)$$

This iteration allows us to map one-to-one from $\phi_g(s)$ into the probability mass function for family size at any generation and estimate values for $N(t)$ in the utility function of equation (13). To specify the process fully we need numerical values for a_k , which we assign below.

Further, using the properties of i.i.d. random sums, it is straightforward to show that if μ is the expected number of daughters in the first generation from a single mother and $\mu < 1$, then the expected family size at generation g is

$$\mathbb{E}(G_g) = \mu^g \quad (19)$$

and

$$\text{var}(G_g) = \sigma^2 \mu^{g-1} \left(\frac{1 - \mu^g}{1 - \mu} \right) \quad (20)$$

where $\text{var}(G_0) = 0$ and $\text{var}(G_1) = \sigma^2$. We can infer that $\mathbb{E}(G_g^2) = \text{var}(G_g) + \mathbb{E}^2(G_g)$.

Steffensen and Christensen proved that the probability that any family line reaches extinction at generation g , x_g , can be computed iteratively by the recursion

$$x_g = a_0 + a_1 x_{g-1} + a_2 x_{g-1}^2 + a_3 x_{g-1}^3 + \dots, \quad (21)$$

where x_{g-1} is the probability of extinction at or before generation $g - 1$.

Likewise, the probability that a family eventually reaches extinction along the female or male line depends on the average number of daughters (sons) born to women (men) in the family. The expected number of children of one gender or the other born to any individual mother or father can be written as

$$\phi'_1(s) |_{s=1} = \mathbb{E}(\xi) = \mu = a_1 + 2a_2 + 3a_3 + \dots \quad (22)$$

In the limit as $g \rightarrow \infty$, $x_g \rightarrow 1$, when $\mu < 1$. So in the case where the expected number of daughters born to mothers is greater than one, the probability of family survival down the female line, $(1 - \rho)$, is non-zero in the limit. In the case where $\mu < 1$, $\bar{F}(\infty) = 0$ and the family will eventually become extinct.

4.1.1 Fertility probabilities

To compute the survival probabilities, we set the probability of family extinction $x_1 = a_0$, where a_0 is the probability that the first female in the family (the establisher of the trust) has no daughters. We assume that the probability that each generation has exactly zero, one, two or more daughters is the same for this particular family as for a representative cohort of mothers, so we can estimate a series for the extinction probability x_g using UK national cohort data on births. We also assume that each subsequent generation has the same constant and known fertility distribution.

Table 1 sets out the probability that a woman born in 1960 in England or Wales has a specified number of live-born children by age 45. These data are taken from Table 10.5 of the UK Office of National Statistics Birth Statistics (FM1 no.34, 2006) and are derived from data on registered birth by year of occurrence and age of mother, but reported in Table 10.5 by mother's year of birth. We use the 1960 cohort of women since they were

the youngest who had reached age 45 at the time of our study and therefore more likely to represent future fertility patterns than older cohorts. We treat the mothers as having no more children after age 45 since the Birth Statistics tables include any births to older mothers in the 45 years data. Average age at first birth for this birth cohort is 27.8, and higher (over 30) for more educated women (see Rendall et al. 2005). In addition we limit the mothers to a maximum of four children each, since the Birth Statistics do not report probabilities for larger families, specifying only the percentage of mothers born in 1960 who at age 45 had ‘four or more’ children.

By assuming that girls and boys are equally likely to be born live (boys are actually slightly more likely to be born than girls, but also suffer higher average mortality for most of life), we also derive the corresponding probability that a woman gives birth to the specified number of girls, where the probability of $K = k$ girls among $R \leq 4$ children is

$$\Pr(K = k) = \binom{R}{k} (0.5)^R. \quad (23)$$

The values in the lowest row of Table 1 are estimates of the probability that a particular family has exactly zero, one, two, three or four or more daughters, that is, a_k in equations (22) and (21). By substituting these values into (22) and checking whether $\mu < 1$, we can infer the overall likelihood of family extinction along the female line. The expected number of daughters to a woman born in 1960 is $0.949 < 1$, which satisfies the condition for eventual family extinction.

[Insert Table 1 here]

Further, by substituting these values into (21), setting the initial probability of extinction at $x_0 = 0$, (the family founder survives in the first generation), and then generating $\{x_g\}_{g=1}^{\infty}$ recursively, we can compute the likelihood that the representative family becomes extinct at or before any particular future generation. The generational survival probability of the single-branch family, $s_g^{(1)} = 1 - x_g$, derived using the probabilities in Table 1, begins at one initially, decreases steeply over the first few generations and converges slowly towards zero, as we can see from the second column in Table 2. The expected family size follows a similar pattern (column four).

[Insert Table 2 here]

The generation g hazard rate, $\lambda_g^{(1)}$, that is the risk of extinction at the current generation conditioning on the family having survived so far, is set out in column three of Table 2. Since the estimated hazard rate is declining with time, we expect that family trustees with rational uncertainty over survival will discount future consumption with decreasing

impatience as the time horizon lengthens.

4.1.2 Approximate distributions

In section 3 above, we proposed that the planning horizon, here limited by family survival, might be exponentially distributed, so that $\bar{F}(t) = \exp(-\lambda t)$, or Pareto distributed so that $\bar{F}(t) = (1 + \beta t)^{-\gamma/\beta}$. The recursively computed survival function in Table 2 represents a discrete analogue to the continuous cumulative survival distribution $\bar{F}(t)$. By fitting both an exponential and a hyperbolic curve to the discrete survival function, we can estimate values for the constant exponential hazard rate λ_{exp} , and the parameters of the hyperbolic function, γ and β .

To find the best fitting continuous distribution function, we calculate 100 generations of discrete survival probabilities s_i , and space each generation 30 years apart (Tremblay and Vezina 2003). We then fit the curve,

$$s_{\text{exp},i} = \exp(-\lambda_{\text{exp}}30i) = \exp(-\lambda_{\text{exp}}t) \quad (24)$$

where $\hat{\lambda}_{\text{exp}}$ is the estimated hazard rate which minimizes the sum of squared errors

$$\min_{\lambda} f(\lambda) = \sum_{i=0}^{100} (s_i - s_{\text{exp},i})^2 = \sum_{i=0}^{100} (s_i - \exp(-\lambda_{\text{exp}}30i))^2. \quad (25)$$

The fitted exponential curve is shown in Figure 1. Here, $\hat{\lambda}_{\text{exp}} = 0.0095$, which is analogous to a constant discrete-time subjective discount factor, of 0.991 per year. In other words, under these assumptions, an expectation of current average rates of family extinction creates mild impatience. However the graph shows that the fit of the function is poor, with the exponential approximation under-predicting and then over-predicting discrete recursive survival probabilities. The sum of scaled squared errors, a guide to the accuracy of the exponential approximation, is $sse_{\text{exp}} = \sum_{i=0}^{100} \frac{(s_i - \hat{s}_{\text{exp},i})^2}{\hat{s}_i} = 3255789$. The mean time to extinction under the estimated exponential distribution is 105.3 years, or 3.51 generations of 30 years.

[Insert Figure 1 here]

A hyperbolic function is a better approximation to the family survival function. Using the same recursively computed discrete survival probabilities, we fit the Pareto function

$$s_{\text{hyp},i} = (1 + \beta(30i))^{-\gamma/\beta}, \quad \beta > 0, \gamma > 0 \quad (26)$$

where the parameters γ and β are estimated by minimizing the sum of squared errors,

$$\min_{\gamma, \beta} f(\gamma, \beta) = \sum_{i=0}^{100} (s_i - s_{hyp,i})^2 = \sum_{i=0}^{100} \left(s_i - [1 + \beta (30i)]^{-\gamma/\beta} \right)^2. \quad (27)$$

Figure 2 shows the fitted curve when the survival function is $\bar{F}(t) = (1 + \hat{\beta}t)^{-\hat{\gamma}/\hat{\beta}}$ and the estimated parameter values are $\hat{\beta} = 0.0112$ and $\hat{\gamma} = 0.0165$. In this case the sum of scaled squared errors is $sse_{hyp} = \sum_{i=0}^{100} \frac{(s_i - \hat{s}_{hyp,i})^2}{\hat{s}_i} = 0.245$, which is much lower than for the fitted exponential curve. The mean of the extinction density is

$$\mathbb{E}(t) = \int_0^{\infty} t\gamma (1 + \beta t)^{-(\gamma/\beta+1)} dt = \left[\frac{1 + t\gamma}{(\beta - \gamma)} (1 + \beta t)^{-\gamma/\beta} \right]_0^{\infty} = \frac{1}{\gamma - \beta} \quad (28)$$

(The limit of the integral in equation (28) as $t \rightarrow \infty$ can be derived by L'Hopital's rule since

$$\lim_{t \rightarrow \infty} \left[\frac{1 + t\gamma}{(\beta - \gamma)} (1 + \beta t)^{-\gamma/\beta} \right] = \lim_{t \rightarrow \infty} \frac{1}{(\beta - \gamma) (1 + \beta t)^{\gamma/\beta-1}}, \quad (29)$$

which goes to zero when $\gamma/\beta > 1$.)

[Insert Figure 2 here]

For the estimated parameter values $\hat{\beta} = 0.0112$ and $\hat{\gamma} = 0.0165$, the expected value of the distribution, or the mean survival of the typical UK family from this cohort is 188.7 years, or 6.3 generations of 30 years. Hence the Pareto distribution predicts a much slower mean extinction time than the exponential distribution.

Had we a longer sample of data on fertility, we could make a statistical comparison between the rival exponential and hyperbolic functions, but that would also entail dealing with some complex issues of testing. Statistically, the exponential distribution is nested inside the hyperbolic distribution, being the special case where $\beta = 0$. This restriction corresponds to a boundary value for the parameter space of beta values for the hyperbolic distribution. Further, the distribution of a test statistic based upon likelihood ratio principles is a weighted sum of chi squared variables with the weights depending upon nuisance parameters. Since we have only one observation on birth patterns we shall content ourselves with maximizing goodness of fit, both on visual grounds, and in terms of sum of squared errors, and continue to work with the assumption that the survival probabilities are known with certainty.

Using the estimated parameters, the formula for the hyperbolic hazard rate is

$$\hat{\lambda}_{hyp} = -\frac{\bar{F}'(t)}{\bar{F}(t)} = \frac{\hat{\gamma}}{1 + \hat{\beta}t} \quad (30)$$

whereas the exponential hazard rate is the constant λ_{exp} . We compare the constant exponential hazard with the hyperbolic hazard in Figure 3, where the 300 years along the horizontal axis corresponds to ten 30-year generations. Over that time the hyperbolic hazard rate declines from around 0.0163 to close to 0.0038, against the constant exponential approximation of 0.0095.

[Insert Figure 3 here]

4.2 Family survival with multiple branches

If, at some arbitrarily chosen generation g_0 , there were found to be k surviving daughters (potential mothers), the future survival density will be determined by the compound probability that all of the k branches of the family reach extinction by a particular future generation, given that k mothers survive at g_0 almost surely. By independence, the probability of extinction at or before generation $g_0 + h$, is $(x_{g_0+h})^k$ where x_{g_0+h} is the probability that a single branch of the family is extinct at or before generation $g_0 + h$ computed using equation (21), and k is the number of branches or originating matriarchs in the family ($x_{g_0} = 0$, for all k). All branches reach extinction in the limit, but the likelihood of survival at any $g_0 + h$ increases as k increases.

Figure 4 graphs the hazard rate at each generation for $k = 1, 2, \dots, 5$, computed using the fertility data from Table 1 and the recursion,

$$s_{g_0+h} = 1 - (x_{g_0+h})^k, \quad x_{g_0+h} = a_0 + a_1x_{g_0+h-1} + a_2x_{g_0+h-1}^2 + a_3x_{g_0+h-1}^3 + a_4x_{g_0+h-1}^4. \quad (31)$$

[Insert Figure 4 here]

When the family has a single branch, the hazard rate is monotonically decreasing, but in the multiple branch family hazard functions are not generally monotone. Figure 4 shows that the empirical hazard rate increases for the first few generations, reaches a maximum, and then declines as the survival c.d.f. approaches zero. For the family with two branches, this maximum is at the second generation, but as the number of branches increases, the maximum moves to a later generation, so that for the five-branch family the maximum occurs at $h = 6$. Consequently a monotonically decreasing discounting

function is not a true reflection of the survival density in multiple-branch families. Table 2 reports the survival probability, hazard rate and expected family membership for a family with three branches in period zero.

If the survival density of the single branch is Pareto as in equation (9), the cumulative extinction density for k branches is

$$F(t|k) = F(t)^k = \left[1 - (1 + \beta t)^{-\gamma/\beta}\right]^k \quad (32)$$

and

$$pdf(t|k) = k \left[1 - (1 + \beta t)^{-\gamma/\beta}\right]^{k-1} \gamma (1 + \beta t)^{-(\gamma/\beta+1)}. \quad (33)$$

The survival function is therefore

$$\bar{F}(t|k) = 1 - F(t|k) = 1 - \left[1 - (1 + \beta t)^{-\gamma/\beta}\right]^k, \quad (34)$$

corresponding to a hazard rate

$$\lambda = -\frac{\bar{F}'(t)}{\bar{F}(t)} = \frac{k\gamma F(t)^{k-1} (1 + \beta t)^{-(\gamma/\beta+1)}}{1 - F(t)^k}. \quad (35)$$

Differentiating (35) with respect to t , and setting the first order condition equal to zero yields an implicit solution for the time where the hazard rate is maximized: choosing t^* to satisfy the equation:.

$$\frac{-\gamma}{\beta} = \frac{\left(1 - F(t^*)^k\right) \left(1 - (1 + \beta t^*)^{-\gamma/\beta}\right)}{\left(1 - F(t^*)^k\right) - k(1 + \beta t^*)^{-\gamma/\beta}}. \quad (36)$$

After this point, the discount function is approximately hyperbolic, but prior to this point, the discount function exhibits increasing hazards.

Applying the multi-branch family survival function will reweight the timing of consumption to favour periods when survival probabilities are high, other thing being equal. Rates of spending will increase towards t^* then decline from that time into the future. In section 5 we report an example of optimal drawdown for a 3-branch family.

Having derived an approximate survival density for a family, we can now apply the analysis of section 3 and section 4 to the trust planning problem.

5 Consumption plans for a family trust

Here we present estimates of the impact of uncertain survival on the optimal drawdown rate of a family trust under the range of scenarios discussed above.

We first solve the problem for a perpetual trust established by a single ancestor with a deed extending to future generations, and derive the optimal path conditioning on information about family fertility at time zero. A second formulation of the problem allows the family to begin with more than one branch and a third approach allows the family to project the number of surviving members at some future point in time and change benefit payments according to the number of existing and expected future members. The second formulation deals with the multiple-branch family still using a unitary utility function. Unlike the single branch case, the multiple branch family does not use a hyperbolic discount function, but a non-monotone function, with slower rates of consumption in the near future and higher, decreasing rates later, consistent with the hazard function described in section 4.2. (We compute the example for a family beginning with three progenitors.) In the third formulation, we report results for the utility function described by equation (13), where the expected future family size influences spending plans.

Consider the discrete-time approximation to the utility maximization problem set out in equation (4). The family trust plans to maximize expected utility over consumption C_t (payments to beneficiaries or disbursements to worthy causes, viewed either as a unitary sum or per family member), by choosing each period a drawdown from uncertain wealth $m_t W_t$, where the gross returns to the trust's investment portfolio are denoted \tilde{Z}_t .

For the first two cases we can work with the same objective function. We represent the probability of family survival at time t by the time-varying parameter, δ_t , where t now takes integer values for years. This parameter can be interpreted as a discrete-time analogue to the continuous cumulative survival density $\bar{F}(t)$ and represents the discount factor at time t . For the single branch family $\delta_t = (1 + \beta t)^{-\gamma/\beta}$ which is monotonically declining, and when there are k progenitors $\delta_t = 1 - [1 - (1 + \beta t)^{-\gamma/\beta}]^k$ which rises and then declines.

Assuming that utility is time-separable and additive, the trust's problem is to maximize expected utility from consumption, L ,

$$L = \mathbb{E} \left(\sum_{t=0}^{\infty} \bar{F}(t) \bar{U}(C_t) \right) = \mathbb{E} \left(\sum_{t=0}^{\infty} \delta_t \bar{U}(C_t) \right), \text{ where } 0 < \delta_t < 1, \quad (37)$$

$$C_t = m_t W_t \quad (38)$$

and

$$W_t = (1 - m_{t-1}) W_{t-1} \tilde{Z}_{t-1}, \quad W_0 = \bar{W}. \quad (39)$$

We can rewrite (39), the difference equation in wealth, as

$$W_t = W_0 \prod_{i=0}^{t-1} (1 - m_i) \tilde{Z}_i. \quad (40)$$

Define the accumulated value of one unit of wealth invested at time $t = 0$ and held until time t , \tilde{V}_{t-1} , as the random non-negative value

$$\tilde{V}_{t-1} := \prod_{i=0}^{t-1} \tilde{Z}_i. \quad (41)$$

If \tilde{Z}_i are also non-negative and independent and identically distributed (i.i.d.), then $(\tilde{V}_{t-1})^\theta$ has a constant mean and $\left[\mathbb{E} \left(\tilde{V}_{t-1} \right)^\theta \right]^{1/t}$ is constant for all t , if the mean exists. Thus equation (40) can be written as:

$$W_t = W_0 \tilde{V}_{t-1} \prod_{i=0}^{t-1} (1 - m_i), \quad (42)$$

and expected utility as

$$L = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta_t \bar{U} \left(m_t W_0 \tilde{V}_{t-1} \prod_{i=0}^{t-1} (1 - m_i) \right) \right]. \quad (43)$$

The first-order condition for optimal drawdown at time t is therefore:

$$\frac{\partial L}{\partial m_t} = \mathbb{E} \left[\delta_t \bar{U}' \left(m_t W_0 \tilde{V}_{t-1} \prod_{i=0}^{t-1} (1 - m_i) \right) \left(W_0 \tilde{V}_{t-1} \prod_{i=0}^{t-1} (1 - m_i) \right) - \sum_{j=1}^{\infty} \delta_{t+j} \bar{U}' \left(m_{t+j} W_0 \tilde{V}_{t+j-1} \prod_{i=0}^{t+j-1} (1 - m_i) \right) \left(m_{t+j} W_0 \tilde{V}_{t+j-1} \prod_{i=0, i \neq t}^{t+j-1} (1 - m_i) \right) \right]. \quad (44)$$

Explicit solutions for the drawdown rate m_t depend on the form of the utility function. For log utility, $\bar{U}(C_t) = \ln(C_t)$, so that,

$$\frac{\partial L_0}{\partial m_t} = \frac{\delta_t}{m_t} - \frac{\sum_{j=1}^{\infty} \delta_{t+j}}{(1 - m_t)} = 0, \quad (45)$$

and

$$m_t = \frac{\delta_t}{\delta_t + \sum_{j=1}^{\infty} \delta_{t+j}} = \frac{\delta_t}{\sum_{j=0}^{\infty} \delta_{t+j}}. \quad (46)$$

Using the condition that at the optimum $\frac{\partial L_0}{\partial m_t} = \frac{\partial L_0}{\partial m_{t+1}} = 0$, we can write the change in the drawdown path as

$$\frac{m_{t+1} - m_t}{m_t} = \frac{\delta_{t+1}}{\delta_t} \left(\frac{\sum_{j=0}^{\infty} \delta_{t+j}}{\sum_{j=1}^{\infty} \delta_{t+j}} \right) - 1. \quad (47)$$

Equation (46) will be constant for a constant δ , so that drawdown depends entirely on the rate of time preference and $\bar{m} = 1 - \delta$. As the term in brackets in (47) approaches one, the proportional change in the optimal drawdown rate varies with the discrete-time hazard rate, $\lambda_{t+1} = -\frac{\delta_{t+1} - \delta_t}{\delta_t}$. For the hyperbolic (Pareto) survival function, this hazard rate is declining over time, so the proportional change in the drawdown also declines. The non-monotonic survival function of the multiple-branch family will be similarly reflected in the consumption path.

For CRRA utility where $\bar{U}(C_t) = \frac{C_t^{1-\alpha}}{1-\alpha}$, and α is the coefficient of relative risk aversion and inverse of the elasticity of intertemporal substitution, an analogous result obtains. Define the risk-adjusted expected return to wealth as $\varphi = \mathbb{E}(\tilde{Z}^{1-\alpha})$ and $\varphi^t = \mathbb{E}(\tilde{V}_{t-1}^{1-\alpha})$, the optimal drawdown at time t when the discount rate is constant is

$$\bar{m} = 1 - (\delta\varphi)^{1/\alpha}. \quad (48)$$

However if the discount rate varies with time, then combining the CRRA utility function with (44) gives

$$m_t = \left\{ \left[\varphi \frac{\delta_{t+1}}{\delta_t} m_{t+1}^{1-\alpha} + \sum_{j=2}^{\infty} \delta_{t+j} m_{t+j}^{1-\alpha} \varphi^j \prod_{i=t+1}^{t+j-1} (1 - m_i)^{1-\alpha} \right]^{1/\alpha} + 1 \right\}^{-1} \quad (49)$$

and the change in the drawdown rate will be

$$\frac{m_{t+1}}{m_t} = \frac{(\delta_{t+1}\varphi)^{1/\alpha} \left[\delta_{t+1} m_{t+1}^{1-\alpha} + \delta_{t+2} m_{t+2}^{1-\alpha} \varphi \left[(1 - m_{t+1})^{1/\alpha} - 1 \right] + \Upsilon \right]^{1/\alpha} + 1}{\delta_t^{1/\alpha} \Upsilon^{1/\alpha} + 1} \quad (50)$$

where $\Upsilon = \delta_{t+2}m_{t+2}^{1-\alpha}\varphi + \sum_{j=3}^{\infty} \delta_{t+j}m_{t+j}^{1-\alpha}\varphi^{j-1} \prod_{i=t+2}^{t+j-1} (1 - m_i)^{1-\alpha}$. Taking logs,

$$\begin{aligned} \ln\left(\frac{m_{t+1}}{m_t}\right) &\approx \frac{1}{\alpha} \ln\left(\frac{\delta_{t+1}}{\delta_t}\right) + \frac{1}{\alpha} \ln(\varphi) \\ &+ \left[\frac{1}{\alpha} \ln\left(\delta_{t+1}m_{t+1}^{1-\alpha} + \delta_{t+2}m_{t+2}^{1-\alpha}\varphi \left[(1 - m_{t+1})^{1/\alpha} - 1\right] + \Upsilon\right) - \frac{1}{\alpha} \ln(\Upsilon) \right]. \end{aligned} \quad (51)$$

So the change in the optimal rate of drawdown will be approximately proportional to $\left(\frac{\delta_{t+1}}{\delta_t}\right)^{1/\alpha} \varphi^{1/\alpha}$, a function of the varying survival probability δ_t , increasing with expected returns and rising as the elasticity of intertemporal substitution rises. Hence the time-varying discounting resulting from uncertain family survival results in subtle but important differences in optimal spending plans when compared with the constant drawdowns under an infinite horizon.

For the third case where we consider the formulation of utility in equation (13), discounting depends on the distribution of future family size at a particular time t , which we denote for the discrete case as the random variable N_t . Since it is necessary to specify a distribution through time rather than at each generation g as we did in section 4.1, we shall make an homogeneity assumption that distributional properties are the same within generations. Drawdown rates are a function of family size in year t , N_t , scaled up by the parameter α , so that the consumption to wealth ratio is

$$m_t = \left\{ \left[\varphi \frac{\mathbb{E}[(N_{t+1})^\alpha]}{\mathbb{E}[(N_t)^\alpha]} m_{t+1}^{1-\alpha} + \sum_{j=2}^{\infty} \mathbb{E}[(N_{t+j})^\alpha] m_{t+j}^{1-\alpha} \varphi^j \prod_{i=t+1}^{t+j-1} (1 - m_i)^{1-\alpha} \right]^{1/\alpha} + 1 \right\}^{-1}. \quad (52)$$

Using the approximation in equation (51) we have

$$\ln\left(\frac{m_{t+1}}{m_t}\right) \approx \frac{1}{\alpha} \ln\left(\frac{\mathbb{E}[(N_{t+1})^\alpha]}{\mathbb{E}[(N_t)^\alpha]}\right) + \frac{1}{\alpha} \ln \varphi, \quad (53)$$

and the growth rate in the drawdown depends on the expected rate of growth of the size of the family scaled by α .

An interesting special case of constant spending arises when $\alpha = 1$. Here the expected family size declines geometrically since $\mathbb{E}[N_{t+1}] = \bar{\mu}^{t+1}$, where $\bar{\mu}$ is the time t (annualised) analogue to the generational mean $\mathbb{E}[G_g] = \mu^g$ and $\mu < 1$ (see equation 19). We estimate $\mu = 0.949$, for our representative UK family, which implies $\bar{\mu} = (\mu)^{1/30} \approx 0.9983$. This constant discount rate means a low constant spending rate $m_t = \bar{m} = 1 - \bar{\mu}$, close to

0.17% p.a.

More generally, consider a family trust whose investment return is 4.75% p.a. in real terms. (This return is close the 15 year historical average for a typical UK trust with a well-diversified portfolio.) Figure 5 sets out a numerical estimate of the first 450 years of the optimal drawdown of a family trust whose survival is modelled by the Pareto distribution estimated in Figure 2 with the curvature parameter, $\alpha = 2$. The thin solid line shows the optimal spending rate using the exponential approximation to single-branch family survival set out in Figure 1, the dotted line is the infinite horizon path where general impatience is set to zero, the dashed line is the optimal path for the single-branch family under Pareto survival and the lighter solid curves are paths for a three-branch family and for the trust concerned with family size. (Changing general impatience away from zero simply moves all curves up the vertical axis by a constant amount.)

[Insert Figure 5 here]

Declining hazard rates create a decreasing shape in the single-branch hyperbolic curve, but the certainty of eventual extinction ensures that both the hyperbolic and exponential drawdown rates are higher than that for an infinitely lived-trust. An ideal spending plan at $\alpha = 2$ for the single-branch unitary trust begins close to 2.1% p.a., and drops toward 1.6%. By contrast, the multiple-branch unitary family trust begins spending close to 1.7% and increases the rate towards 1.8% before tapering off. The effect of a lower elasticity of intertemporal substitution ($\alpha = 2$) on the trust which distributes by number of members is to produce a conservative early drawdown (close to 1% per year), which is later increased above the infinite horizon path, before declining very gradually in later generations.

6 Conclusion

Recent studies (Sozou, 1998; Dasgupta and Maskin, 2005; Bommier 2006) have shown that decreasing impatience can be a rational response to horizon uncertainty. This is a very common feature of long-term planning problems, and family trusts are just one example of the many bodies that must consider stochastic ‘survival’. Indeed we all have to plan for uncertain lifetimes. By contrast with family survival from a single progenitor, which we have modelled using a Pareto distribution with hyperbolically declining hazard rates and risk neutral survival preference, individual mortality (at least later in life) is better fitted by the increasing hazards typical of a Gompertz function. An increasing

hazard rate suggests rationally increasing impatience, perhaps motivating the elderly aunt who says ‘but I’ll be dead by then’ as a reason for not planning as far out as next Christmas.

In the standard intertemporal consumption model with i.i.d. returns, time-varying hazards mean time-varying optimal drawdown rates, a result that goes against the customary advice to trusts and endowments to spend at a constant rate. Our results have interesting implications for family foundation trustees. Estimates for UK families along the female line, assuming that current fertility patterns persist into the future, signal eventual extinction for a typical family. Faced with the resulting hyperbolic survival function, it seems hard for a trustee behaving in the interest of the family dynasty to justify a policy of constant consumption. The ideal plan for a perpetual trust which treats consumption among members as perfectly substitutable and does not review its survival function, is to spend more rapidly in the near future, and steadily but more slowly as the trust ages. The best plan for a multiple-branch family is likely to be a non-monotone path, exhibiting increasing then decreasing impatience through time. When the individual consumption of trust beneficiaries is key, rather than simple family survival, risk aversion over wealth interacts with expected family growth to determine the optimal path. Constant drawdown is optimal for trusts with a power utility parameter at unity but drawdown is low early, slowly increasing and then later declining for trusts which are less willing to transfer consumption through time.

Our results could be applied to the survival of financial institutions such as banks, mutual funds or hedge funds, or to more general macroeconomic questions such as the estimation of a social discount rate. We have also set aside the issue of risk aversion over family mortality. Both questions we leave to future work.

Acknowledgements:

We thank participants at the School of Economics and Finance, Queensland University of Technology, Seminar, and the Econometrics Society Australasian Meetings 2008, two anonymous referees and the editor for helpful comments. Thorp acknowledges support from ARC DP0877219.

References

- Ainslie GW (1974) Impulse control in pigeons. *Journal of the Experimental Analysis of behaviour* 21(3):485-489.
- Albertsen K (1995) The Extinction of Families. *International Statistical Review* 63(2):234-239.
- Albertsen K, Kristensen E (1995) A Previously Unpublished Solution to Erlang's "Problem of Extinction of Families" by Carl Marius Christensen. *International Statistical Review* 63(2):242-245.
- Becker GS (1974) A Theory of Social Interactions. *Journal of Political Economy* 82(6):1063-1093.
- Becker GS (1981) Altruism in the Family and Selfishness in the Market Place. *Economica* 48(189): 1-15.
- Becker GS, Barro, RJ (1988) A Reformulation of the Economic Theory of Fertility. *Quarterly Journal of Economics* 103(1):1-25.
- Becker GS, Tomes, N (1986) Human Capital and the Rise and Fall of Families. *Journal of Labor Economics* 4(3) Part 2: The Family and the Distribution of Economic Rewards, S1-S39.
- Blanchard OJ (1985) Debts, Deficits and Finite Horizons. *Journal of Political Economy* 93(2):223-247.
- Bommier A (2006) Uncertain Lifetime and Intertemporal Choice: Risk Aversion as a Rationale for Time Discounting. *International Economic Review* 47(4):1223-1246.
- Burke J (1976) *Osborn's Concise Law Dictionary*. 6th edition, Sweet & Maxwell, London.
- Cigno A (1993) Intergenerational Transfers without Altruism: Family, Market and State. *European Journal of Political Economy* 9(4): 505-518.
- Cigno A (2006) A Constitutional Theory of the Family. *Journal of Population Economics* 19:259-283.
- Cigno A (2007) *Children and Pensions*. MIT Press, Cambridge MA.
- Dasgupta P, Maskin E (2005) Uncertainty and Hyperbolic Discounting. *American Economic Review* 95(4):1290-1299.

- Harris C, Laibson D (2001) Dynamic Choices of Hyperbolic Consumers. *Econometrica* 69(4):935-958.
- Harris TE (1963) *The Theory of Branching Processes*. Springer-Verlag, Berlin/New York.
- Hubbard RG, Skinner J, Zeldes SP (1995) Precautionary Saving and Social Insurance. *Journal of Political Economy* 103(2):360-399.
- Hull DM (1998) A Reconsideration of Galton's Problem (using a two-sex population). *Theoretical Population Biology* 54(2):105-116.
- Ingersoll JE (1987) *Theory of Financial Decision Making*. Rowman and Littlefield, Savage.
- Kendall DG (1966) Branching Processes since 1873. *Journal of the London Mathematical Society* 41:385-406.
- Keyfitz N (1968) Introduction to the Mathematics of Population. Addison-Wesley, Reading M.A.
- Korn R, Korn E (2001) Option Pricing and Portfolio Optimisation. American Mathematical Society, Providence.
- Laibson DI (1998) Life-cycle Consumption and Hyperbolic Discount Functions. *European Economic Review* 42(3-5):861-871.
- Laibson DI, Repetto A, Tobacman J (1998) Self-control and Saving for Retirement. *Brookings Papers on Economic Activity* 1998(1):91-196.
- Light A, McGarry K (2004) Why Parents Play Favorites: Explanations for Unequal Bequests. *American Economic Review* 94(5):1669-1681.
- Loewenstein G, Thaler RH (1989) Anomalies: Intertemporal choice. *Journal of Economic Perspectives* 3(4):181-193.
- Lotka AJ (1931) The Extinction of Families, I. *Journal of the Washington Academy of Sciences* 21:377-380.
- Manser M, Brown M, (1980) Marriage and Household Decision-making: a Bargaining Analysis. *International Economic Review* 21(1):31-44.
- McElroy MB, Horney MJ (1981) Nash-bargained Household Decisions: Toward a Generalization of the Theory of Demand. *International Economic Review* 22(2): 333-349.

OECD (2007) Can Policies Boost Birth Rates? *Organisation for Economic Co-operation and Development Policy Brief*, November.

Office for National Statistics (2006) Birth Statistics. Series FMI no. 34, London.

Phelps ES, Pollack RA (1968) On Second-best National Saving and Game-equilibrium Growth. *Review of Economic Studies* 35(2):185-199.

Pitacco E (2004) Survival Models in a Dynamic Context: a Survey. *Insurance: Mathematics and Economics* 35(2):279-298.

Ramsey FP (1928) A Mathematical Theory of Saving. *Economic Journal* 38(152):543-559.

Rendall M, Couet C, Lappegard T, Robert-Bobee I, Ronsen M, Smallwood S (2005) First Births by Age and Education in Britain, France and Norway. *Population Trends* 121, UK Office of National Statistics: 27-34.

Satchell SE, Thorp S (2007) Scenario Analysis with Recursive Utility: Dynamic Consumption Paths for Charitable Endowments. Working paper, University of Technology, Sydney.

Sozou PD (1998) On Hyperbolic Discount and Uncertain Hazard Rates. *Proceedings of the Royal Society of London (Series B)* 265:2015-2020.

Steffensen JF (1995) On the Probability that the Offspring Dies Out. *International Statistical Review* 63(2):239-242.

Stern N (2006) *Stern Review on the Economics of Climate Change*. Cambridge University Press, Cambridge.

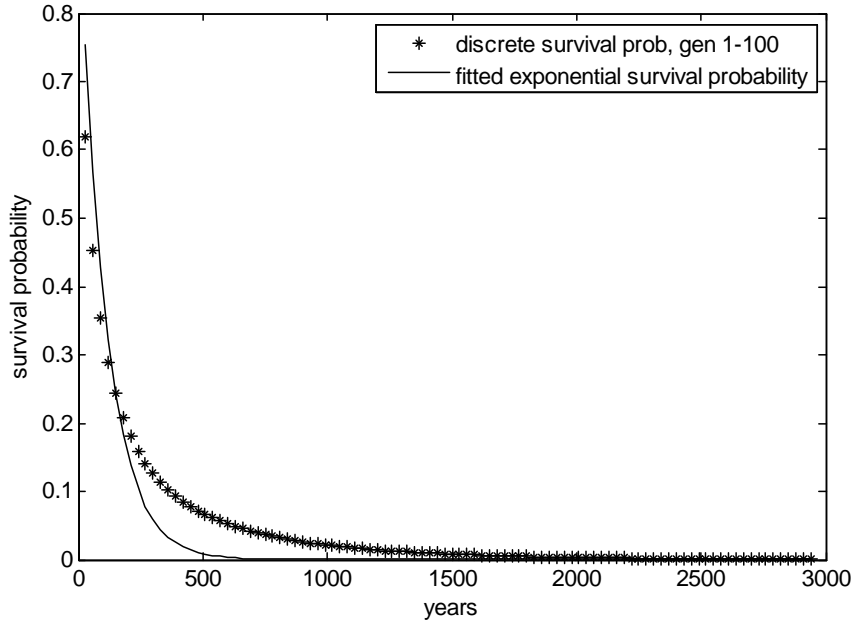
Taylor HM, Karlin S (1998) *An Introduction to Stochastic Modelling*. 3rd edition, Academic Press, San Diego.

Tremblay M, Vezina H (2003) New Estimates of Intergenerational Time Intervals for the Calculation of Age of Origins of Mutations. *American Journal of Human Genetics* 66(2):651-658.

Xu Z (2007) A Survey on Intra-household Models and Evidence. American Institutes for Research Working Paper, Washington DC.

Yaari ME (1965) Uncertain Lifetime, Life Insurance and the Theory of the Consumer. *Review of Economic Studies* 32(2):137-150.

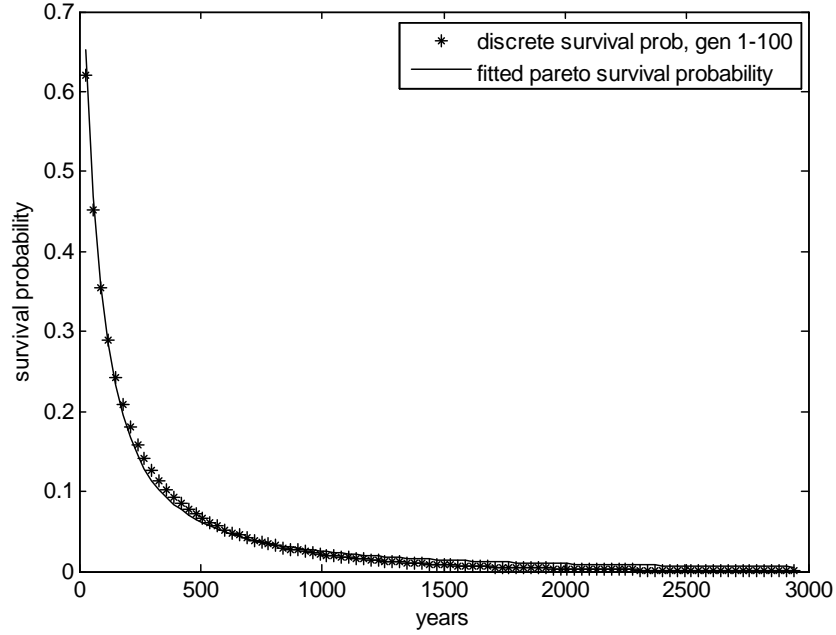
Figure 1: Fitted exponential survival function for a single branch family.^a



^aFigure shows a graph of the fitted exponential function $\hat{s}_{\text{exp},i} = \exp(-\hat{\lambda}_{\text{exp}}t)$ where s_i is the probability of family survival at generation $i = t/30$ and t are years. The generation $i = t/30$ survival probabilities are calculated recursively from equation (21) along the female line for the 1960 birth cohort of English and Welsh women, assuming that the likelihood of the birth of 0-4 girls exactly is constant over time and homogeneous across the population. (See Table 1 and Table 2.) Function is fitted by `fminsearch` in Matlab using a simplex method for non-linear optimization to minimize the sum of squared errors


$$\min_{\lambda} f(\lambda) = \sum_{i=0}^{100} (s_i - \exp(-\lambda_{\text{exp}}30i))^2.$$

Figure 2: Fitted hyperbolic survival function for a single branch family.^a




^aFigure shows a graph of the fitted hyperbolic function $\hat{s}_{hyp,i} = [1 + \hat{\beta} (30i)]^{-\hat{\gamma}/\hat{\beta}}$ where s_i is the probability of family survival at generation $i = t/30$ and t are years. The generation $i = t/30$ survival probabilities are calculated recursively from equation (21) along the female line for the 1960 birth cohort of English and Welsh women, assuming that the likelihood of the birth of 0-4 girls exactly is constant over time and homogeneous across the population. (See Table 1 and Table 2.) Function is fitted by `fminsearch` in Matlab using a simplex method for non-linear optimization to minimize the sum of squared errors $\min_{\gamma, \beta} f(\gamma, \beta) = \sum_{i=0}^{100} (s_i - [1 + \beta (30i)]^{-\gamma/\beta})^2$.

Figure: 3: Estimated hazard rate for single-branch family.^a




^aFigure shows the estimated exponential hazard rate $\hat{\lambda}_{\text{exp}} = 0.0095$ and the estimated hyperbolic hazard rate $\hat{\lambda}_{\text{hyp}} = -\frac{\bar{F}'(t)}{\bar{F}(t)} = \frac{\hat{\gamma}}{1+\hat{\beta}t}$ where $\hat{\beta} = 0.0112$ and $\hat{\gamma} = 0.0165$. See Figures 1 and 2 and the text for estimation details.

Figure 4: Estimated hazard rate for multiple-branch family.^a



^aFigure shows the probability of family extinction and survival down the female line where there are k original mothers at generation g_0 and the probability of survival at each future generation is given by, $s_{g_0+h} = 1 - x_{g_0+h}^k$, $x_{g_0+h} = a_0 + a_1x_{g_0+h-1} + a_2x_{g_0+h-1}^2 + a_3x_{g_0+h-1}^3 + a_4x_{g_0+h-1}^4$. The a_k are the probabilities that a mother has exactly k daughters. Values of a_k are taken from the last row of Table 1, the estimated distribution of daughters to the cohort of mothers born in England and Wales in 1960. The hazard rate is the probability of extinction between generation $g_0 + h$ and $g_0 + h - 1$, conditioning on having survived to time $g_0 + h - 1$, which is computed by $\lambda_{g_0+h} = -\frac{s_{g_0+h} - s_{g_0+h-1}}{s_{g_0+h-1}}$

Figure 5: Optimal drawdown with survival uncertainty.^a



^aFigure shows the estimated optimal spending rates with and without uncertainty over family membership and family survival for trusts with power utility preferences $U(C_t) = \frac{1}{1-\alpha}C_t^{1-\alpha}$, (exponential, single-branch, triple-branch and infinite horizon curves) or $U(C_t) = \frac{1}{1-\alpha}C_t^{1-\alpha}N_t^\alpha$, $\alpha = 2$, (family size $N(t)$ curve) and investment returns close to 4.75% p.a. Estimates of the risk-scaled investment returns are bootstrapped from historical portfolio returns to a typical investment trust (See Satchell and Thorp 2007). The hyperbolic survival probability of the single-branch family is given by the distribution function $\bar{F}(t) = (1 + \beta t)^{-\gamma/\beta}$, $\hat{\beta} = 0.0112$ and $\hat{\gamma} = 0.0165$, of the triple-branch family as $\bar{F}(t) = 1 - [1 - (1 + \beta t)^{-\gamma/\beta}]^3$ and the exponential survival probability is $\bar{F}(t) = \exp(-\hat{\lambda}_{\text{exp}}t)$, $\hat{\lambda}_{\text{exp}} = 0.0095$. The curvature parameter-scaled expectation of future family size in the $E(N_t^\alpha)$ is derived from the recursion equation (18). Numerical optimization is via the fminimax routine in Matlab.

Table 1:
 Estimated distribution of women of child-bearing age by number of children.^{a,b}

Population proportion of women born 1960 having number of children (live births) by age 45				
0	1	2	3	4 or more
0.18	0.13	0.38	0.20	0.10
Estimated probability of women born 1960 having number of daughters (live births) by age 45				
0	1	2	3	4
$a_0 = 0.380$	$a_1 = 0.355$	$a_2 = 0.208$	$a_3 = 0.050$	$a_4 = 0.007$

^aSource: Table 10.5, Birth Statistics, Series FM1 no.34, Office for National Statistics, London, UK.

^bTop panel shows population data on births for women born 1960 in England and Wales. Lower panel shows inferred probability of daughters by assuming that girls and boys are equally likely to be born so if the probability that a woman from this cohort has 2 children is 0.38, the probability that one will be a girl is 0.5×0.38 and that 2 will be girls is 0.25×0.38 etc. We use these as estimates of the probabilities a_k that a progenitor (mother) has exactly zero, one, two, three or four daughters. We assume that the family consists of female members and the number of daughters born to each (female) family member is a random draw from this distribution at every generation.

Table 2:
Estimated probability of family survival and expected family size.^a

Number of Generations	Single branch family			Triple branch family		
	$s_g^{(1)} = 1 - x_g$	$\lambda_g^{(1)}$	$\mathbb{E}(G_g^{(1)})$	$s_g^{(3)} = 1 - (x_g)^3$	$\lambda_g^{(3)}$	$\mathbb{E}(G_g^{(3)})$
0	1.000		1.000	1.000		3.000
1	0.620	0.380	0.949	0.945	0.055	2.847
2	0.452	0.271	0.901	0.836	0.116	2.702
3	0.354	0.217	0.855	0.731	0.126	2.564
4	0.289	0.183	0.811	0.641	0.123	2.433
5	0.243	0.160	0.770	0.566	0.117	2.309
6	0.208	0.144	0.730	0.503	0.111	2.191
7	0.181	0.131	0.693	0.450	0.105	2.080
8	0.159	0.121	0.658	0.405	0.100	1.974
9	0.141	0.113	0.624	0.366	0.096	1.873
10	0.126	0.106	0.592	0.333	0.092	1.777
11	0.114	0.100	0.562	0.303	0.088	1.687
12	0.103	0.095	0.534	0.277	0.085	1.601
13	0.093	0.092	0.506	0.255	0.082	1.519
14	0.085	0.088	0.481	0.234	0.080	1.442
15	0.078	0.085	0.456	0.216	0.078	1.368
16	0.072	0.081	0.433	0.200	0.076	1.298
17	0.066	0.080	0.411	0.185	0.074	1.232
18	0.061	0.077	0.390	0.172	0.072	1.169
19	0.056	0.076	0.370	0.159	0.071	1.110
20	0.052	0.073	0.351	0.148	0.069	1.053
25	0.037	0.066	0.285	0.105	0.064	0.854
30	0.026	0.064	0.219	0.077	0.060	0.657
35	0.019	0.059	0.169	0.057	0.058	0.506
40	0.014	0.053	0.130	0.042	0.056	0.389
45	0.011	0.053	0.100	0.032	0.055	0.300
50	0.008	0.047	0.077	0.024	0.054	0.024
150	≈ 0.000		≈ 0.000	≈ 0.000		≈ 0.000

^aTable shows the probability of family survival and expected family size down the female line where the probability of extinction for a single-branch family at generation g is given by

$x_g = a_0 + a_1x_{g-1} + a_2x_{g-1}^2 + a_3x_{g-1}^3 + a_4x_{g-1}^4$, for a three-branch family is $(x_g)^3$, and a_k is the probability that a mother has exactly k daughters. Values of a_k are taken from the last row of Table 1, the estimated distribution of daughters to the cohort of mothers born in England and Wales in 1960. The recursion begins with $x_1 = a_0 = 0.38$ and continues with fixed a_k . Hazard rates are the probability of extinction between generation $g - 1$ and g , conditional on having survived to time g , which is

$$\text{computed by } \lambda_g^{(j)} = -\frac{s_g^{(j)} - s_{g-1}^{(j)}}{s_{g-1}^{(j)}} : s_g^{(1)} = 1 - x_g, \text{ and } s_g^{(3)} = 1 - (x_g)^3.$$