


ALPHA SOURCES

NOVEMBER, 2018



THE LIFE CYCLE HYPOTHESIS

How individuals and economies bargain with the future

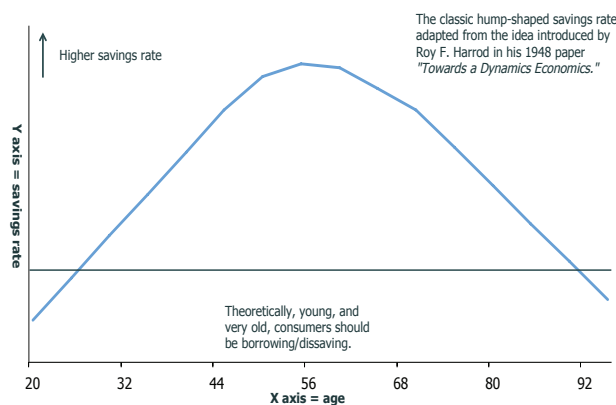
Depending on the wording, a search on Google Scholar for papers and research on the link between macroeconomics and demographics yields anywhere between 20,000 and 100,000 results. This is an unhelpful start for someone looking to explore and understand the field. This essay aims to rectify this issue by tracing the origins of the life cycle hypothesis (LCH)—ground zero for linking macroeconomics and demographics—through a close inspection of the 1950s literature that gave birth to the theory. It is motivated by the idea that anyone who wishes to explore this topic needs to have a firm grasp of the original material. A lot can be said for getting the basics right, and this essay is an ode to that idea.

WHAT IS THE LIFE CYCLE HYPOTHESIS?

The LCH is based on the notion that economists analyse the consumption and saving decisions of economic agents as a function of their transition from young, to middle-aged, to old and, eventually, retired consumers. If taken seriously, this leads to two general conclusions. First, it implies that consumers are forward-looking in their decision-making. In theory, this assumption is banal; but in practice, the introduction of uncertainty to the analysis of consumer behaviour poses a big challenge to economic theory, even to this day. Second, the LCH suggests that decisions regarding consumption and savings are independent of the level of, and change in, current income.

The LCH is formulated as a microeconomic theory, but it has a logical macroeconomic counterpart. If consumers base their consumption decisions on how old they are, and how long they expect to live, it follows that the aggregate age structure can tell us something about the behaviour of the economy as a whole. In a 2005 lecture in Rome, Angus Deaton, an American economist, gave one of the most eloquent descriptions of the theory. The lecture starts with the simple observation that young people save today so that they can spend later in life, when they either cannot, or prefer not to, work. This idea gives rise to the classic hump-shaped savings schedule, which links the rate of saving to the age of the individual.

If aggregated to the economy as a whole, this implies that the wealth of a nation is passed down through generations. The young and middle-aged have relatively little wealth, and will save to buy the assets that the elderly need to sell in order to finance consumption in old age. In an economy with a growing (working-age) population, the savings of the young will be higher than the dissavings of the old; growth will be positive and wealth will rise. But what happens if the working-age population doesn't grow? Furthermore, what is the consequence of long-term growth falling or even reaching zero, as would appear to be the case in rising number of economies? The need for an inquiry into such quandaries is just one of the reasons why the LCH still matters today.



BEYOND THE NEOCLASSICAL HERITAGE

Winston Churchill allegedly said that "History is written by the victors"; an effective metaphor for the history of the LCH. Pick up any advanced macroeconomic textbook and Samuelson's 1958 paper, *An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money*, and Diamond's 1965 contribution, *National Debt in a Neoclassical Growth Model*, tend to be cited as the origins of the LCH by way of an overlapping generations (OLG) model. Franco Modigliani, who won a Nobel Prize for fathering the LCH, also gets due attention, although not always to the extent that he perhaps deserves.

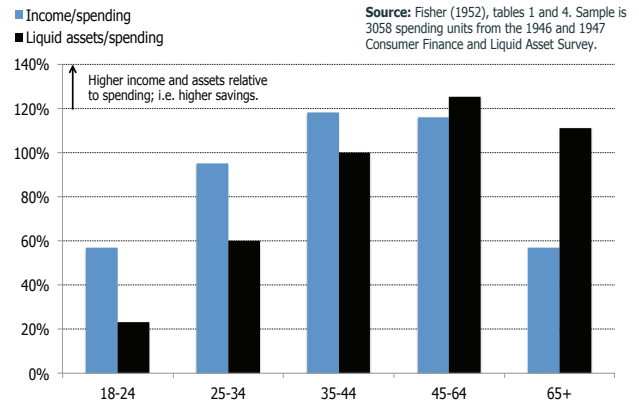
▶▶ In most cases, the neoclassical imperative wins; Samuelson and Diamond provide the origins for a standard two-period representative agent model used as part of a neoclassical growth model. This lineage is consistent with the fact that, by the end of 1950s, the neoclassical economic approach was well on its way to total domination. Samuelson's and Diamond's papers are important, and I will devote attention to them later, but they are by no means the original sources of the LCH, nor do they necessarily convey the best description of such an idea.

WHAT DO THE DATA SAY?

The early empirical literature on the relationship between ageing and consumers' consumption decisions has mostly been forgotten. This is a great shame, since you will be hard-pressed to find a better synthesis of the relationship between macroeconomics and demographics than these contributions. Our journey starts with Fisher's study of spending and saving patterns in the U.S., with data from the 1935-36 Consumer Purchases Study, the 1946 Liquid Assets Survey and 1947 Survey of Consumer Finances.

Fisher's opening paragraph is a remarkable piece of writing. It starts by noting that the median age of the U.S. population is rising, and that such a trend is likely to continue. In the beginning of the 1950s, Fisher couldn't have known that the U.S. was in the midst of the post-war baby boom, though in fairness, she acknowledges the early evidence of this process. That said, her hypothesis that the median age, as well as the proportion of elderly people, eventually would rise has been proven right in the U.S., and, more spectacularly, in the majority of countries in the Organisation for Economic Co-operation and Development. She effectively summarises the entire LCH research programme when she writes: "The central question of this study is whether changes in the age composition of the population influence the national income and proportion saved".

The study itself is, as far as I know, the forerunner of an increasingly large body of literature in which consumer survey and cohort expenditure data are used to conduct cross-sectional and, in some cases, time series analyses. These studies offers incredibly important evidence, but they suffer from methodological problems. In an ideal world, researchers would be able to observe an entire cohort from cradle to grave across the entire cross section. But this is beyond even modern-day survey and data collections.



Fisher's study uses data from the 1947 Survey of Consumer Finance to analyse a sample of households—3058 in total—immediately after the Second World War. Generalising these results to the population is problematic. We can't be sure that middle-aged consumers in a given year are representative of middle-aged consumers in the future, any more than we can be sure that the snapshot of behaviour in a given year is an unbiased sample of behaviour over the life cycle and life course as a whole.

In the case of Fisher's study, it is tempting to argue that consumer behaviour immediately after an event as important as the Second World War is almost certainly not an unbiased indicator. Finally, survey design and data collection methods differ across time and countries, which means that trying to generalise across countries and time is difficult.

Without a line of algebra, Fisher derives the idea of the LHC from the distinction between two different cycles; an income and an expenditure cycle. The extent to which these two differ over the life cycle—and they almost certainly do—the rate of savings is determined by subtracting one from the other. The key question is simple: what happens to consumption and saving as households age?

The first table presented by Fisher is in the form of a savings rate—the ratio between income and spending—which peaks between the ages of 35 and 64 before declining for spending units—the subject of analysis in Fisher's sample—for ages above 65. The highest level of average savings is found in the age group 45-64. These results fit the idea that households dissave to fund spending in old age, and that they save the most just before retirement.

This conclusion, however, is qualified in two ways. First, the decline in income is not matched by a decline in assets. The ratio of liquid assets to spending

▶▶ peaks for the age group 45-64, and dips only slightly for units aged above 65. The *flow* of savings declines as households age, but the stock barely falls.

Second, dissaving in old age is correlated with income to a much higher extent than in other age groups. In short, the rate of saving for older households tends to increase with income, which, by extension, suggests that the rate of dissaving is higher for older households with low income.

Indeed, Fisher's results provide strong evidence that dissaving in old age is heavily concentrated among households with low incomes. By contrast, the data suggest that elderly households with high incomes tend to have savings rates matching those found in middle-aged groups. In other words, dissaving in old age is not an optimal strategy for those households that don't need to do so.

This has profound consequences if generalised to the economy as a whole; namely, that changes in age structure are associated with shifts in inequality. Fisher makes the following point: "It [the result] suggests that, other things being equal, the ageing of a population over a long period may be associated with substantial changes in income structure".

Another conclusion is that ageing societies are not necessarily characterised by dissaving; at least not until very late in their demographic transition. This is controversial, especially in international macroeconomics where many standard models of capital flows assume that ageing economies need inflows to sustain external deficits. Empirically speaking, ageing economies—Japan, Germany, etc.—seem to be characterised by excess savings and the propensity to run structural and large external surpluses. To put this in macroeconomic terms, it appears that investment demand declines faster than the desired rate, and stock, of savings as populations age. Even if microeconomics suggests that ageing eventually leads to outright dissaving, this is not necessarily the optimal response for the economy as a whole.

Lydall's 1955 paper, *The Life Cycle in Income, Saving, and Asset Ownership*, is methodologically similar to that of Fisher. In other words, it is inductive. It starts with observations and attempts to formulate a theory by the end of the analysis. Lydall is at pains to point this out at the offset of the paper: "In the social sciences, empirical study of society must usually take precedence over deductive reasoning; in econometrics, the mathematician must be the ally, not the master, of the economist". It is easy to see

why the inherently deductive neoclassical tradition has opted to look beyond the studies of Fisher and Lydall in writing the history of the LCH.

The study by Lydall is based on data from a series of U.K. national surveys carried out in the 1950s by the Institute of Statistics at Oxford, which focused on personal income and saving. More specifically, the paper uses data from 1953, in which 2,100 respondents were included. The comparison with Fisher's study is clear, and Lydall gives due attention to his colleague in the introduction. The challenge of equivalence, however, is obvious to the reader from the outset. It is not clear that results from a 1947 survey in the U.S. can be readily compared with the information from a 1953 survey in the U.K., not least because the sample designs also differ considerably.

As it turns out, Lydall's results chime with Fisher's in most key areas. Savings and income appear to follow a classic hump-shaped form, though the sources of income shift across the life cycle from labour to a mix of government transfers and asset income. In addition, relatively high-income households have considerably higher savings rates than their middle- to low-income counterparts. Finally, households in Lydall's sample also dissave into old age—the savings rate peaks at about 40—but their net worth increases steadily until just before retirement, declining only marginally afterwards.

Charles Zwick's small 1958 study, *Demographic Variation: Its Impact on Consumer Behavior*, is an apt way to end this part of the essay. It takes its point of departure in the results of Fisher and Lydall, and asks whether income and price elasticity of consumption vary with age. The study is consistent with the findings of Fisher and Lydall.

In a sample of 151 families in Medford Massachusetts—over a period of 32 weeks between 1952 and 1953—age is found to be the only statistically significant variable explaining income and price elasticity. Specifically, the results indicate that older consumers—aged 40 and above—have a relatively low, short-term price elasticity, but a high income elasticity. It is the opposite for young households.

The idea behind such a phenomenon is that young consumers exhibit a low income elasticity to food, but have a high income-elasticity towards durable goods. As households age, income elasticity towards durables declines in line with a reduced need to replace existing products. This, in turn, drives a rise in income elasticity with respect to food. Conversely,

▶▶ price elasticity is high for young consumers as their habits are less fixed than those of old consumers.

Robert Zwick's study is a rather dubious attempt to prove that the ageing process exerts a meaningful influence over consumers' consumption decisions. But it is a good example of someone's difficulty in attempting to replicate previous empirical studies; an exercise that modern empirical macroeconomics struggles with, too.

THE BIRTH OF A THEORY

My second review of Franco Modigliani's seminal work in the 1950s—I did the first in my master's thesis—and his Nobel Prize lecture in 1985 leads to a controversial conclusion: at its inception, the key objective of the LCH was *not* to construct a link between demographics and consumer behaviour, but rather to mount an assault on the Keynesian analysis of consumption and saving. Establishing a link between an ageing population and the way it chooses to consume and/or save was simply a means to an end.

This interpretation is easiest to defend by noting that Modigliani was developing his theory at the same time as Milton Friedman was putting the finishing touches on the ground-breaking "*A theory of the consumption function*", which provided the foundation for the permanent income hypothesis (PIH). The subplot here is simple; the LCH and PIH were both formulated at a point in time when economists were thinking about ways to differentiate themselves from the Keynesian analysis of consumption and saving.

It is fair to say that Modigliani and Friedman approached the critique of Keynes from two separate angles. The PIH is an attempt to differentiate the Keynesian dictum that consumers spend out of current income as a function of their marginal propensity to consume. Modigliani is interested in the determinants of saving, and why the Keynesian idea that too much saving leads to inadequate aggregate demand—the paradox of thrift—doesn't necessarily hold over time. In truth, however, the LCH and PIH are complimentary theories. They both try to establish a framework capable of analysing consumption and saving as a forward-looking decision made, to some extent, under uncertainty.

The LCH has been the subject of countless of papers and articles; Modigliani's Nobel Prize lecture and Angus Deaton's 2005 paper, Franco Modigliani and the Life Cycle Theory of Consumption, are the best secondary sources. The opening salvo, however, was

made by the 1954 paper by Brumberg and Modigliani, *Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data*.

In contrast to Fisher and Lydall, Brumberg and Modigliani's paper develops a theoretical model, and the authors make their position clear when they say: "By now, the amount of empirical facts that has been collected [about consumption] is truly impressive; if anything, we seem to be in imminent danger of being smothered under them. What is, however, still conspicuously missing is a general analytical framework which will link together these facts".

A simple question lies at the heart of Brumberg and Modigliani's research: why do people save? In asking this, they are able to identify four reasons. First, the desire to leave a bequest is a powerful, and simple, motivation. Second, saving arises because the flow of current and expected income does not necessarily coincide with the preferred path of consumption. Most consumers prefer a relatively stable consumption path over their lifespan, but their income path is almost surely anything but. This motivation gives rise to the intuitive idea that the working-age years should be spent saving for retirement.

The two aforementioned motivations drive savings even under perfect certainty, but uncertainty adds a further two. The precautionary motivation is a catch-all concept for the kind of savings consumers make for emergencies or unexpected events. A particularly interesting source of uncertainty is life expectancy, and the idea that consumers do not know when their lifespan will end. No one can expect to optimise their savings path by throwing in the towel the day that their savings, excluding desired bequests, reach zero. The final motivation is logical, but trivial. Due to the uncertainty surrounding one's lifespan, acquiring assets, houses, cars, etc. requires consumers to put up equity, which functions as invested savings.

Brumberg and Modigliani get to the core of their message when they present the idea that the "marginal propensity to consume" (MPC) is independent of income, but dependent on age. Specifically, the MPC is an increasing function of age, and, by extension, the "marginal propensity to save" is the opposite. It is an *inverse* function of an increase in age. This, however, only holds in a model with perfect foresight, where individuals know the length of their lifespan.

It is ironic that this definition of the life cycle has survived as the main version of Modigliani's theory. Almost everything that has since been written on the

▶▶ LCH has either presented it as a stylised version of the world, which doesn't hold in reality, or attempted to outright disprove it. Indeed, the early studies by Fisher and Lydall are cases in point in that they both suggest that decisions about savings in old age differ considerably from the standard model.

Angus Deaton suggests that Modigliani was annoyed by this interpretation: "I remember Modigliani being driven to something close to apoplexy by the (once upon a time often-heard) suggestion that the LCH is a theory for bachelors, that it ignores both the presence of children early in the life-cycle, and bequest motives at its end. This 'bachelor' claim can only reflect a failure to read the work".

I am sympathetic to Modigliani's objections, but the shortcomings of the basic model still lurk ominously over this discussion like a dark family secret that no one dares to talk about. This is because many theoretical studies that use some variant of the OLG model still argue that the transition into old age will lead to aggregate dissaving, eventually. In an open economy, this translates into the prediction of capital imports, and external deficits, in rapidly ageing economies. Most empirical studies along the lines of Fisher and Lydall suggest that this is unlikely, but they cannot refute the theoretical claim that dissaving will set in at some point.

The extent to which dissaving is an inevitable equilibrium of a rapidly ageing economy is primarily challenged by countries such as Japan and Germany. They appear to be fighting the onset of dissaving to the bitter end, however this may look. In other words, rapidly ageing economies today tend to be characterised by significant excess savings and external surpluses. This is a problem.

It is one thing to argue that a path towards dissaving is an uneven one, but if ageing economies act in opposition to the tenets of the LCH, the theory needs a rethink. Earlier in this essay I referred to a description of the LCH by Angus Deaton, who emphasises the ability of the LCH to explain growth and saving in an economy with positive population growth and a growing working-age population. It is fair to say, however, that the LCH struggles to explain the behaviour in economies where the working-age population is shrinking; a pressing challenge given the present trend of rapidly ageing populations. Even if this is true, it shouldn't be seen as an indictment of the theory, but a call for reflection; perhaps the theory simply needs a rethink?

THE NEOCLASSICAL STORY

Paul Samuelson's contributions to economics in the middle of 1940s are still considered some of the most important ideas of economic theory to come out of the twentieth century. It seems apt then that Samuelson's 1958 paper, alongside Modigliani's early work, has survived, and since been identified, as the origins of the LCH. In similar fashion to Modigliani, Samuelson does not start with the idea that shifts in age structure can explain macroeconomic phenomena.

Instead, he sets out to provide a "general equilibrium solution to the determination of the time-shape of interest rates". Samuelson converges on the idea of a life cycle, when he establishes the idea that "men want to consume less than they produce during their working years, so that they can consume something in the years in which they produce nothing". This is the type of language that makes people frown at economics, so let me try to explain it in plain English: to develop a theory for value of resources through time, which is what an interest rate is, you need to imagine the ability to make a bargain with the future.

Samuelson proves this by examining the life of Robison Crusoe, who can only trade with nature. If the interest rate is zero, whatever Crusoe saves as he works will be available to him when he retires. If the capital saved—for example, coconuts or other crops—is inherently productive, it will earn a positive interest, and compounding will boost Crusoe's wealth. In the extreme situation where capital perishes immediately, no trade with nature is possible, and Crusoe dies the minute he is no longer able to work.

From here, Samuelson transforms the economy into a dynamic system by assuming three generations, in which older workers (with wealth) cede capital to younger workers (with no wealth) as a "bribe" so that they are taken care of when they retire and can no longer work. The laws of general equilibrium then take over to set up a simple economy.

For the economic modeller to sleep calmly, he must find the values of interest rates, both for today and tomorrow, which satisfy the constraint that the discounted value of consumption by three generations equals the discounted value of their production. This rule of thumb is at the heart of all neoclassical research on the LCH; the general equilibrium is a binding constraint. Anyone with graduate-level economics under their belts knows this. The imperative to find a tractable mathematical solution condemns the model to a deceptively simple expression of the world.

▶▶ The LCH is most often cast in the form of a two- or three-period OLG model in which a representative consumer optimises his lifetime path of consumption and saving. The end result—see the appendices [here](#) for an example—is usually an Euler equation, which can either be estimated directly or linearised via a Taylor expansion. By this point, however, the mathematical component has usually become at least as important as economic intuition in gaining the final result. This is also why I won't go into great detail about Diamond's 1965 paper, *National Debt in a Neoclassical Growth Model*, since it takes its point of departure in neoclassical growth theory, which is a topic for another essay.

AN IDEA TO CHERISH, AND DEVELOP

The LCH has all the hallmarks of a good economic theory. It is intuitive, simple and profound. In its youth, it helped economists to move the theory of consumption and saving beyond the 1930s analysis set out by Keynes. It has matured with grace, providing key insights into the consumption and saving decisions of individuals as they age. It also introduces the investigation into how demographic transitions drive the flow of savings within and between economies. In a world where populations are ageing at an unprecedented pace, it is as important as ever.

The LCH matters for policymakers in that it helps them design public pension systems and rules for retirement. The theory can help governments to put into place incentives to make sure that households have adequate savings for retirement. For investors, the LCH matters because it provides insight, at least in theory, into the direction of capital flows in a world with free capital mobility and countries in different stages of their demographic transitions.

As we venture deeper into the first half of the twenty-first century, though, the LCH needs a tune-up. Considerable work has been done to explain why the model, in its simplest form, fails to capture consumption and saving decisions in old age. Consumers ought to draw their savings down as they age, but precautionary motivations and the desire to leave bequests mean that they don't, at least not to the extent predicted by a standard model. This is well understood, but many applications of the LCH still assume a stage of dissaving in old age, which is a problem. Even if the rate of savings declines with age, consumers tend to hold on to their assets well into retirement. This appears to hold even before cor-

recting for the fact that many elderly consumers stay longer in the labour force via part-time employment.

This argument extends to the analysis of economies as a whole, especially in a world of free capital mobility. Economies such as Germany, Japan and Italy are different in terms of culture and economic models, but they all seem to be converging on the same growth path. Excess savings compared to low domestic investment demand necessitates the need to run external surpluses. The onset of dissaving might well be a terminal state, but it is not an optimal one. Economies don't retire, and notwithstanding war or natural catastrophe, they don't die. This matters when trying to aggregate the microeconomic conclusions of the LCH to the economy as a whole. In short, the simple model tends to get it wrong, especially when analysing economies that are further ahead in their demographic transitions.

This puts the analysis in a quandary. If rapidly ageing economies are somehow breaking the laws of economics by running external surpluses, what would they look like without them? Initial conditions matter. Could Germany, Italy or Japan finance persistent external deficits? Who would buy their assets, and at what price? Could these economies afford to import capital and foreign savings to finance domestic investment in excess of what they can fund themselves? I doubt it. Instead, ageing economies—and especially those who have suffered recent balance-of-payments stress—appear to have been dragged into a cul-de-sac where only one model works.

At this point, I suspect many of my fellow economists will accuse me of attacking a straw man. The LCH doesn't have to perfectly account for global capital flows, or the presence of imbalances in global current accounts, to be a good theory. Both of these can't be reduced to a simple function of ageing populations anyway. It is becoming increasingly clear that excess private savings are accumulating in the corporate sector, though this is not necessarily an issue for the LCH, if households own the firms that are doing the saving. It will take a separate essay to link these stories to the validity, or lack thereof, of the LCH and onset of rapid population ageing.

When all is said and done, though, I think the main points of this essay stand. The LCH is one of the core economic theories of consumption and saving, and in a world with a rapidly ageing population, and still-free capital mobility, it should be taken seriously by investors, researchers and policymakers alike.

▶▶ LIST OF REFERENCES

Harrod, Roy (1948) - *Towards a Dynamic Economics*, Macmillan; (Dec. 1948).

Fisher, Janet A (1952) - *Income, Spending and Saving Patterns of Consumer Units in Different Age Groups*, Studies in Income and Wealth Volume 15.

Lydall, Harold (1955) - *The Life Cycle in Income, Saving and Asset Ownership*, *Econometrica*, Vol. 23 No. 2 (Apr 1955), pp. 131-150.

Modigliani, Franco and Brumberg, Richard (1954) - *Utility Analysis and the Consumption Function: An Interpretation of the Cross-Section Data*, Originally published in Post-Keynesian economics. - London : Allen and Unwin. - 1955, p. 388-436.

Friedman, Milton (1957) - *A Theory of the Consumption Function*, *The Review of Economics and Statistics*, Princeton University Press.

Zwick, Charles (1958) - *Demographic Variation: Its Impact on Consumer Behavior*, *The Review of Economics and Statistics* Vol. 39, No. 4 (Nov., 1957), pp. 451-456.

Samuelson, Paul A (1958) - *An Exact Consumption-Loan Model of Interest with or without the Spacial Contrivance of Money*, *The Journal of Political Economy* Vol. 66, No. 6 (Dec., 1958), pp. 467-482.

Diamond, Peter A (1965) - *National Debt in a Neoclassical Growth Model*, *The American Economic Review* Vol. 55, Issue 5 (Dec., 1965), pp. 1126-1150.

Modigliani, Franco (1985) - *Life Cycle, Individual Thrift and the Wealth of Nations*, Nobel Prize lecture, December 9 1985. The PDF can be retrieved [here](#).

Deaon, Angus (2005) - *Franco Modigliani and the Life Cycle Theory of Consumption*, Presented at the Convegno Internazionale Franco Modigliani, Accademia Nazionale dei Lincei, Rome, February 17th–18th, 2005. The PDF can be retrieved [here](#).

Vistesén, Claus (2010) - *Demographics and the Anatomy of International Capital Flows*, Master's Thesis at Copenhagen Business School. The PDF can be retrieved [here](#).

ACKNOWLEDGEMENTS

I would like to thank Rosie Tate and Peter Clayburn from the *Proofreading Company* for providing valuable edits and comments on this essay.