

On the value of not taking expert advice

Guy Thomas

E-mail: R.G.Thomas@kent.ac.uk

SUMMARY

Conventional wisdom commonly exhorts non-experts to take expert advice when dealing with specialist fields. This works well in relation to the physical or biological world, because theories of these worlds are generally neutral: popular acceptance of a theory does not change the phenomena it describes. In contrast, theories of social phenomena such as finance are often reflexive: popular acceptance of a theory does change the phenomena it describes. Reflexive theories can be either self-fulfilling or self-negating. Advice based on self-negating theories is not likely to be useful. Expert advice is therefore less useful in fields such as investment, which are dominated by self-negating theories.

1. Introduction

In many specialist fields of knowledge – law, medicine, investment – it is a truism that a non-expert who needs to interact with the field should seek expert advice. This truism glosses over the point that expert advice is less useful in some fields than in others. The value of expert advice depends on the *type* of theory on which the expert's knowledge is based. Advice based on theories which are *self-negating* is often not useful, and so in fields which are dominated by self-negating theories, the value of taking expert advice is relatively low. Investment is one such field. It is not obvious that people who follow expert advice in investment do much better than people who ignore it.

In this article, the term 'theory' is used broadly and encompasses any scheme for understanding the world and making predictions or deciding what to do. 'Taking advice' implies a process of consultation whereby the non-expert – the client – presents his idiosyncratic problem to an expert, and then relies on advice based on the expert's theories. An expert in a particular field is any person who meets some threshold of professional competence, which is usually demonstrated by the possession of relevant qualifications. Thus the term expert encompasses any suitably qualified person; it is not restricted to persons with exceptional knowledge or authority. By way of example, all solicitors in public practice are experts on the law in their particular fields of practice, and all private client stockbrokers are experts on investment; the term expert is not limited to lawyers who write law textbooks, or investors with exceptional track records. The antithesis of taking advice, 'not taking advice,' means that the non-expert does not follow the paradigm of consultation and reliance on an expert's advice. It does not imply capricious contrarianism, nor does it mean that the non-expert engages in no research and deliberation. Indeed, a person who is 'not taking advice' may well be motivated to engage in *more* self-directed research and deliberation. This is one of the possible advantages of not taking advice.

2. Neutral, self-fulfilling, and self-negating theories

An engine, not a camera

Theories in physical and biological sciences are generally *neutral*, the sense that popular acceptance of the theory has no effect on the phenomena which the theory represents. For example, the phenomenon of evolution and the truth of theories which describe it are unaffected by popular belief or scepticism. But in social sciences, theories which purport to explain or

describe phenomena can often alter those phenomena, and so increase or reduce the truth of the theories. Theories in social science do not just *describe* or *predict* outcomes, they also *create* outcomes – that is, theories are *reflexive*. The point is succinctly made by the title of a book by the English sociologist Donald Mackenzie, theories about financial markets are ‘*An engine, not a camera*’ⁱ.

Reflexive theories can be divided into two types: self-fulfilling or self-negating. *Self-fulfilling theories* are those where the popular acceptance of the theory increases its truth. *Self-negating theories* are those for which the truth engine runs in reverse: popular acceptance of the theory reduces its truth. Theories in investment are often reflexive: sometimes self-fulfilling, but more often self-negating. Neutral theories are also found in finance. Some examples of each type of theory in investment are given below.

Neutral theories

The arithmetic of compounding is an example of neutral theory in finance. Another example is the theory of diversification: adding more shares to a portfolio reduces the variance of its return above and below the return on the market index. These theories do not depend on popular acceptance, and nor are they negated by it.

Self-fulfilling theories

Rational self-interest The idea of rational self-interest as a cornerstone of economics is probably a self-fulfilling theory. If people are taught that other people are always selfish, they may be more likely to be selfish themselves. There is some evidence that university students of economics are more selfish than peer students in other subjects, and that this effect becomes stronger with each extra year’s study of economics.ⁱⁱ

Bank runs A theory which predicts there will be a bank run in particular circumstances is likely to have self-fulfilling properties. It is rational for bank depositors who appreciate the theory to withdraw their deposits before other depositors do so, thus causing the run which the theory predicts.

Black-Scholes option pricing The Black-Scholes theory as published in 1973 relied on unrealistic assumptions, and initially actual option prices in the market did not conform well to the theory. But within a few years, popular acceptance of the theory helped to create the world it described, in at least two ways. First, market makers and other market participants increasingly referred to the theory in setting prices. Second, the theory provided lobbying fodder and legitimacy for regulatory changes which made some of its previously unrealistic assumptions more realistic. For example, the theory assumes an investor can borrow unlimited amounts at the riskless rate of interest to hedge his holdings of option with shares (or vice versa). In practice, in 1973, regulation restricted borrowing and high interest was charged, reflecting the perception that options were a gambling instrument. But because Black-Scholes conceptualised options as a hedging rather than gambling instrument, that legitimised regulatory relaxations on borrowing for option holders, thus making the theory’s assumption of unlimited borrowing at the riskless rate closer to the truth.ⁱⁱⁱ

The benefits of vaccination Outside the field of finance, vaccination against infectious diseases such as measles and mumps provides a particularly compelling example of a self-fulfilling theory. The truth of a theory extolling the benefits of vaccination for any individual depends on a sufficient fraction of the whole population believing and acting upon it.

Self-negating theories: running the truth engine in reverse

Market efficiency: the Grossman-Stiglitz paradox If more investors come to believe that markets are efficient, fewer perceive incentives to conduct fundamental research, and the supply and demand forces which push prices towards efficient levels are weakened. Thus markets can be efficient only if a sufficiently large fraction of investors believe that markets are *not* efficient, so that they are motivated to do fundamental research. This conundrum, the Grossman-Stiglitz paradox, implies that market efficiency is to some extent a self-negating theory.^{iv}

Portfolio insurance The theory of portfolio insurance was developed in the United States in the early 1980s. The theory suggested that the downside risk of an equity portfolio could be limited by selling progressively larger volumes of index futures contracts as share prices fell. By 1987 over \$60bn of US equity portfolios were protected by some form of portfolio insurance. But in the 1987 crash, portfolio insurance failed to limit losses as expected, with index futures sales often being executed far below the intended levels. Because portfolio insurance requires mechanical sales as the index falls with no reference to fundamentals, the widespread use of portfolio insurance programmes probably worsened the crash. The protection which the theory of portfolio insurance was supposed to provide was negated by widespread use of the theory.

academic economists: people for whom a theory's truth represents one of its less important properties, well below tractability and ideological conformance

The small company effect In 1981 the economist Rolf Banz published an article demonstrating that shares in small companies in the United States had provided higher risk-adjusted returns over the previous 40 years than shares in large companies.^v Subsequent studies showed a similar effect in historical data for other markets. However, in the following three decades this 'small company premium' has reduced, or even reversed in many countries – probably because people have acted upon the theory, redirecting their investment research and funds towards small company shares, so that they became less neglected and undervalued.

'Low-risk' investments Any theory which claims a particular financial instrument offers unusually low risk relative to its returns will tend to be self-negating: as the theory becomes popular, demand for the investment bids up its price, which reduces its expected returns and increases its downside risk. For example, in the mid-2000s mortgage-backed securities (MBS) based on pools of house mortgages in the United States were priced and marketed on theories which implicitly assumed that prolonged nationwide falls in house prices were impossible. Widespread acceptance and use of these theories facilitated the decline in lending discipline which negated the theories, turning MBS into high-risk investments.

Careers advice Outside the field of finance, some types of careers advice may be an example of self-negating theory. Advice to students to enter fields expected to be subject to skills shortages and high demand may be self-negating, because the increase in applicants to these fields will reduce or even reverse the predicted skills shortages.

Although a distinction has been drawn above between self-fulfilling and self-negating theories, in some cases the distinction may not be definitive. A reflexive theory may be initially self-fulfilling, as an increasing fraction of the investment community comes to accept and act upon it; the same theory may later become self-negating, when too large a fraction of the community has already acted upon it. This complex feedback between reflexive theories and events – initially self-fulfilling, later self-negating, and then possibly repeating again – can give rise to a boom-

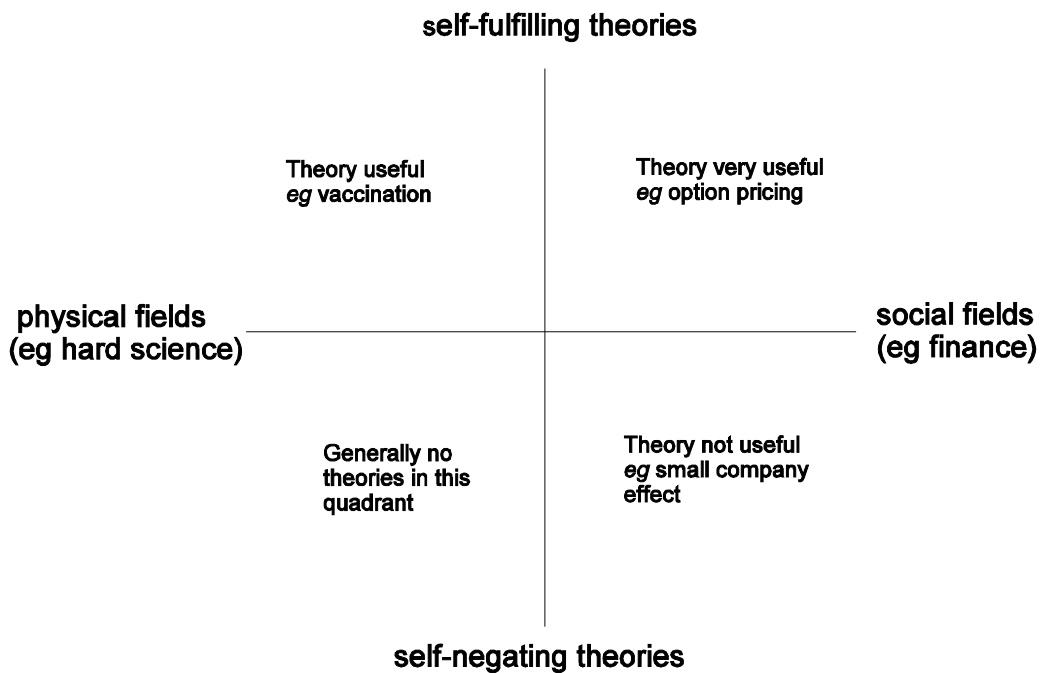
and-bust cycle, as described in several books by George Soros.^{vi} These books are derided by most academic economists: people for whom a theory's truth represents one of its less important properties, well below tractability and ideological conformance.

Value of advice depends on type of theory

Relying on expert advice is straightforward when the expertise is based on neutral theories. Expert advice is therefore usually useful in fields where neutral theories predominate, such as science and engineering. But in social fields such as finance, where many theories are reflexive, the position is more complicated. If a theory is reflexive and self-fulfilling, advice based on it can be very useful. But if a theory is reflexive and self-negating, expert advice based on it is often not useful. People who have experience in scientific fields where theories are mainly neutral tend to expect too much from expert advice when they apply themselves to social fields where theories are often self-negating. The relative prevalence of reflexive theories in physical and social fields, and the implications for the usefulness of theory, are summarised by the quadrants in Figure 1.

Figure 1: The theory quadrants

Did Insert / picture / from file / Quadra2.eps on big pc to get this)



3. Constraints and biases affecting the providers of advice

Apart from the prevalence of self-negating theories in investment, there are often other disadvantages to relying on expert advice. Expert thinking is prone to constraints and biases which non-experts who do their own thinking and make their own decisions can often avoid.

Reputational and legal constraints

Explicit constraints on experts include the reputation, legal liability, and professional codes of practice. Constraints such as these guide the expert towards fashionable orthodoxy: to be safe the expert needs to give the conventional advice which any reasonably competent and up-to-date expert practitioner would be expected to give. These constraints place both upper and lower limits on quality of advice: the threat of legal liability and professional censure inhibits negligence, but it also inhibits original thinking and effective but unconventional solutions. Another explicit constraint is any fee quote or estimate which the expert has given (or any tacit expectation by the expert of the maximum the client may be willing to pay), combined with the expert's valuation of his own time. This places an effective limit on the time which can be spent considering a particular matter, irrespective of any unexpected complexity or importance which emerges.

The explicit constraints of reputation, legal liability and professional codes may be particularly detrimental in fields dominated by self-negating theories. Advice which recognizes the self-negating nature of a profession's popularly accepted theories will generally be unconventional advice, which carries a risk of professional and legal sanctions for the adviser if it later turns out to be wrong, and no benefits for the adviser if it turns out to be right. For an adviser, it is more attractive to give conventional advice based on popular theories than to be right.

The agency problem

In many cases advice is shaped partly by the interests of the adviser, rather than the advised. This is the 'agency problem.' It is particularly prevalent where the adviser's fee is contingent on a particular action being taken. It is probably unwise to ask a barber if one needs a haircut, or a stockbroker if one should buy shares. Taking a particular risk may be optimal for the client, but not optimal for the adviser; and vice versa for other risks. The influence of experts' financial self-interest on advice can be overcome where there is strong value system or professional ethics in place to counteract it. This may sometime be true in medicine, but it is usually not true in finance.

Advice is often supplied with risk warnings. These tend to be more for the protection of the adviser from criticism or legal liability, rather than useful decision inputs for the advised. 'Protective' risk warnings are a distraction for the client, and often divert attention away from more important risks which are difficult to quantify or describe, or for which legal precedents idiosyncratically absolve the adviser of responsibility.

Over-confidence

Advisers often claim to know more than they really know. In many cases, they actually *believe* they know more than they really know. This 'epistemic arrogance' has been repeatedly demonstrated in experiments where people are asked to estimate a range for quantity with a given

confidence level: the error rate of the estimated intervals is always much higher than that consistent with the requested confidence level.^{vii} In other cases, the adviser knows that he does not know, but believes (perhaps correctly) that pretending he does know is more impressive to clients. Taking expert advice can buttress an illusion of knowledge. By not taking advice, it is more likely you understand that you do not know.

Strategic complexity

Advisers often prefer complexity, or at least the illusion of complexity, in their advice for reasons unrelated to the interests of the advised. First, by giving the impression that expert advice is essential, spurious complexity promotes demand for the adviser's services. Second, it facilitates higher fees: complexity can be either a direct justification for overt high charges, or an indirect means of hiding high effective charges: complex charging structures with features designed to confuse customers are very common in finance.^{viii} Third, in situations subject to a high degree of randomness, complicated models or methods give an illusion of control, which may have a calming effect on other stakeholders. Fourth, and most generally, complexity can be used by advisers rather as a priest uses incense – not to clarify, but to impart an air of mysterious authority. There is a good deal of evidence that complexity in analysis beyond a certain point does not on average improve the quality of decisions, and that relatively simple decision rules perform at least as well as more complicated ones in a wide range of contexts.^{ix}

Disciplinary bias

Many problems are best approached with an open mind and as few *a priori* mental constraints as possible. Experts start with the mental constraints of their discipline, of which they exaggerate the importance, whilst downgrading the importance of other types of expertise. The expert would argue that his proper role is to give advice strictly limited to his own area of expertise, and that the non-expert should integrate advice from experts in various fields and form his own judgments. But this is usually unrealistic: in many situations, a non-expert will consult just a single expert, choosing the field of expertise which seems most relevant (for example, a lawyer for a perceived legal problem). Often the expert's judgment then simply substitutes for the client's judgment, so that the problem is considered solely within the constraints of the expert's professional discipline.

4. Advantages of not taking advice

There are many advantages of not taking advice, some of which are applicable in almost any field. First, it often saves time and money: it may be quicker and cheaper to acquire some expertise and deal with a problem yourself, rather than attempt to convey all its factual nuances and your own preferences to an expert. Second, it increases your autonomy and sense of agency: you become the master of your own destiny. Third, it promotes robustness and transparency: when you think for yourself, you understand problems and your own decisions better. Fourth, it is more flexible: you can make decisions at the last possible moment, with the benefit of any late new information, and change your mind easily. Fifth, it is often advantageous to use perspectives from more than one domain. Experts tend to stick to their own domains, and may be unaware or dismissive of the insights of other domains.

One might summarise these advantages by saying that not taking advice promotes the hacker mentality. The word hacker is used here in the laudatory sense in which it acts as a shibboleth amongst good programmers: hackers are not criminals, they are intellectually playful outsiders who can figure out unfamiliar systems, invent unconventional solutions, and get things done for themselves. The best hackers are world

<p>The best hackers are world hackers, willing to hack not just with computers, but with anything</p>

hackers, willing to hack not just with computers, but with anything. The actions of a hacker, like those of a person who takes no advice, are often denigrated by establishment experts as inappropriate, unauthorised, unprofessional, or unqualified. Disparagement in such vague yet *ad hominem* terms is usually a validation of success: when experts are unable to make any meaningful criticisms of substance, they are reduced to whining about the hacker's status as an outsider.

There are some further advantages of not taking advice, which are not universal but specific to particular fields. Not taking advice generally makes you more open to contrarian and unorthodox decisions, which are often intrinsically valuable in investment. Not taking advice also fosters originality, and so may be helpful in fields where some element of creativity is needed.

Not taking advice has been made easier and more attractive in the past two decades by greater accessibility of specialist knowledge to anyone. Two decades ago, specialist knowledge in books and journals was difficult for the non-expert to access. Experts were *de facto* gatekeepers to all specialist information, and experts in each locality had an effective monopoly on their field of expertise. Nowadays, the abstracts for expert articles in all fields are generally free to anyone, and full articles can be conveniently accessed at modest charge. Two decades ago, it was difficult for most people to make contact with people facing similar circumstances (whether medical, legal or in some other field) and swap experiences and ideas. Nowadays, online communities and bulletin boards make this easy. These profound changes make it much easier for the intelligent outsider to hack specialist fields. The changes are routinely downplayed or derided by experts: for example, doctors are often derogatory and patronising about patient online communities, and similarly stockbrokers about investor communities, and so on. But experts are not disinterested evaluators of these changes, because they tend to undermine the esteem and livelihood of experts.

Not giving advice

As well as not taking advice, good decision-making may also be helped by not *giving* advice. This is because giving advice implies a public commitment to a particular position, but avoiding such commitment makes it easier to change your mind. In many situations, it is helpful to keep multiple alternative hypotheses in your mind in parallel, and make decisions only at the last possible moment. Giving advice diverts attention and effort from *making* decisions to *communicating* decisions. The neat narrative required for clear communication is often antithetical to the 'parallel alternative hypotheses' mode of thinking which fosters good decision-making.

Caveat: situations where it helps to take advice

Although the main point of this article is the advantages of not taking advice, we should briefly acknowledge that there are some scenarios where advice may be worth taking. First, there is the point made by Keynes: *'Worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally.'*^x Because of the spurious merit which influential others often attach to taking expert advice, doing this – or more precisely, *appearing* to do this – can sometimes protect you against later criticism. For example, when you have already decided on a particular course of action, but anticipate that it might later be questioned by a court or regulator, it may be advantageous to procure some supportive expert advice by judicious 'opinion-shopping'. Failure or error by a person acting on expert advice (Keynes' 'conventional failure') is often seen as less blameworthy than that by a person who makes his own decisions. For some people, this trick also works on their internal scorecard: if they have acted 'responsibly'

by taking advice, rather than thinking for themselves, they may feel less regret about a subsequent bad outcome. This seems perverse, but it works for some people.

Operating without advisers in contexts where others expect you to have them can provoke unease or even antagonism amongst counterparties and their advisers. But this can have its own advantages: an absence of advisers can increase the perceived unpredictability of your actions, or make you appear more naïve than you really are. In situations involving negotiation, these can sometimes be strategic advantages. However in fields such as litigation, where advisers are not mandatory but nevertheless an almost universal norm, the advantages just mentioned are probably outweighed by the strong prejudice of arbiters such as the judiciary against unadvised parties.

It is interesting that the field in which taking advice seems still seems most advisable, the law, is a notoriously reactionary profession, with a culture of oleaginous deference towards the judiciary, an elite of higher average age than the elites of most professions. Two decades ago, the paradigm of an omniscient expert representing and advising a docile client also prevailed in other fields of expertise, such as medicine and investment. But in these fields, with less deferential cultures and elites which are themselves younger, technology has changed the relationship between expert and client. The most reactionary field with the most elderly elite has been the slowest to change.

In some contentious domains, the prevalence of *one-way hash* arguments may mean it is better to defer to expert opinion than to attempt to evaluate all the arguments for yourself. A one-way hash is an incorrect and tendentious argument which is designed to be intuitive, plausible and appealing to non-experts. Understanding the flaw in the one-way hash argument is more difficult, because it requires a deeper level of understanding. The term one-way hash is an allusion to public-key cryptography, which relies on the fact that it easy to multiply together two large prime numbers, but much more difficult to perform the reverse process of finding the common factors of their product. One-way hash arguments are prevalent in popular discussion of theories which happen to be the subject of political debate, such as global warming or evolution. In these domains, relying on the consensus of expert opinion may be a better strategy than attempting to examine all competing theories yourself, because many theories from outside the expert consensus are not developed in good faith. As in all the other situations described in this article, understanding the process and context in which theories are developed may be a more useful first step than understanding the theories themselves. It just so happens that in domains such as evolution and global warming, understanding the process suggests that experts are probably less conflicted than the non-experts. But the opposite is usually true in investment.

A *one-way hash* is an incorrect and tendentious argument which is designed to be intuitive, plausible and appealing to non-experts

5. How should one deal with fields dominated by self-negating theories?

This article has already noted that because self-negating theories are common in finance, the value of expert advice in finance is often low. Despite this, regulation and conventional wisdom in many countries promote the paradigm of non-experts seeking advice from certified experts in finance. In part, this may be because of cognitive regulatory capture: financial regulators are usually drawn largely from the industry which they regulate, and often promote the interests of the industry rather than customers in their regulation. It may also be because of a lack of appreciation of the low value of advice in a field dominated by self-negating theories.

How then should you deal with such fields? In principle, you could seek exceptional experts who understand and make allowance for the self-negating nature of many of the theories on which

their advice is based. Some sophisticated experts actually do this, but it seems doubtful that the run-of-the-mill qualified practitioners who are characterised as experts in this article will do so. The more practical solution for most people, who do not have access to exceptional experts, is to develop your own awareness of the reflexive nature of theory, and the limited value of advice. You should think about not just the substance of a theory, but also the context and purpose for which it was developed.

Another possible response is to be sceptical of complicated financial theories, which may ultimately turn out to be self-negating, and instead focus on robustness rather than optimisation – or in more colloquial terms, not try to be too clever. ‘Fast and frugal’ decision rules have been shown to work quite well in a variety of contexts. For example, in selecting an investment portfolio it has been shown that a ‘1/n’ allocation which allocates equal amounts to every available asset is often superior to the more sophisticated optimization methods which an expert adviser might typically follow. This is because the gain from optimal weightings relative to the naïve ‘1/n’ equal weightings is typically smaller than the loss in accuracy from running an optimisation model with more inputs, which are estimated with error rather than known precisely.^{xi}

You can also focus and act upon those theories in a field which do not appear to be self-negating. For example, in finance, the theory that increasing the diversification of a portfolio brings the expected return on the portfolio closer to the expected return on the whole market is not self-negating: everyone can act on this theory without it being made less true. In careers advice, guidance about which fields are expected to suffer from skills shortages may be self-negating, but guidance about matching individual attributes to careers is not self-negating. You should therefore pay more attention to the latter type of guidance.

ⁱSee Mackenzie, D. (2006) *An engine, not a camera: how financial models shape markets*. MIT Press.

ⁱⁱFor a survey see Frank, R.H., Gilovich, T. and Regan, D.T. (1993) ‘Does Studying Economics Inhibit Cooperation?’, *Journal of Economic Perspectives*, 7: 159–171

ⁱⁱⁱSee Mackenzie, D. (2003) ‘An Equation and its Worlds: Bricolage, Exemplars, Disunity and Performativity in Financial Economics’, *Social Studies of Science*, 33:831-868. Mackenzie uses the term *performative* where I use self-fulfilling, and *counter-performative* where I use self-negating.

^{iv}Grossman, S.J. and Stiglitz, J.E. (1980) ‘On the impossibility of informationally efficient markets,’ *American Economic Review*, 70: 393-408.

^vBanz, R. (1981) ‘The Relationship Between Return and Market Value of Common Stock’, *Journal of Financial Economics*, 9: 3-18.

^{vi}For example, Soros, G. (1985) *The Alchemy of Finance*.

^{vii}Kahneman, D., Slovic, P. and Tversky, A. (1982) *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.

^{viii}For a discussion of this see Carlin, B.I. (2009) ‘Strategic Price Complexity in Retail Financial Markets,’ *Journal of Financial Economics*, 91: 278-287.

^{ix}Gigerenzer, G., Todd, P.M. (1999) *Simple Heuristics That Make Us Smart*. Oxford University Press.

^xKeynes, J.M. (1936) *The General Theory of Employment, Interest and Money*. Macmillan.

^{xi}For an investigation of the ‘1/n’ rule, see DeMiguel, V., Garlappi, L. and Uppal, R. (2009) ‘Optimal Versus Naïve Diversification: How Inefficient is the 1/n Portfolio Strategy?’, *Review of Financial Studies*, 22: 1915-1953. For fast and frugal rules generally, see Gigerenzer, G., Todd, P.M. (1999) *Simple Heuristics That Make Us Smart*. Oxford University Press, Oxford.