



Information, Uncertainty & Espionage

Peter J Phillips¹  · Gabriela Pohl¹

Accepted: 11 June 2022 / Published online: 28 June 2022
© The Author(s) 2022

Abstract

Intelligence scholars are drawing on behavioural decision theory to improve decision-making under risk and uncertainty in intelligence and counterintelligence. Such an undertaking is essentially lacking without the Austrian school's concepts of knowledge, discovery, (entrepreneurial) judgement, ignorance, rational calculation and, more generally, its analysis of human action in the face of true uncertainty. Decision theory, both orthodox and behavioural, depicts decision rather narrowly as a prioritisation task undertaken within a delineated problem space where the probabilities “sum to one”. From such a perspective, certain perennial challenges in intelligence and counterintelligence appear resolvable when in fact they are not, at least not when approached from the usual direction. We explain how Austrian concepts can complement efforts to improve intelligence decision-making. We conclude that the future strategic value of intelligence analysis is located beyond information acquisition, however fast and however vast. Intelligence agencies have no price signals to help them determine how much intelligence to produce. And governments have no price signals to moderate their appetites for the intelligence product. Ultimately, those agencies that recognise the implications of intelligence agencies as non-price institutions and adapt their decision-making processes may find that they have the upper hand over their rivals.

Keywords Austrian School · Knowledge · Discovery · Ignorance · Uncertainty · Intelligence · Counterintelligence · Behavioural decision theory · Non-price institutions

✉ Peter J Phillips
phillips@usq.edu.au

✉ Gabriela Pohl
gabriela.pohl@usq.edu.au

¹ Faculty of Business, Education, Law and Arts, University of Southern Queensland, 4350 Toowoomba, QLD, Australia

1 Introduction

Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don't know we don't know. And if one looks throughout the history of our country and other free countries, it is the latter category that tends to be the difficult one—Donald Rumsfeld.¹

It might be said that intelligence agencies create the market for intelligence. After all, the acquisition and utilisation of information against one's enemies involves multiple transactions, many of which have a monetary dimension. The agencies themselves speak of 'intelligence product' and its 'consumers'. The institutional structures of intelligence networks are not dissimilar to the evolved routines that guide market participants and the day-to-day operations of all sophisticated organisations. And the ebb and flow of resource allocations to and within intelligence agencies are made in response to signals of various kinds. In many respects then, it seems quite reasonable to speak of a market for intelligence and, from there, it seems but a small step to apply the straightforward economics of supply and demand to the intelligence product, its producers, and its consumers. But if we were to use price theory to evaluate non-price institutions, we would run up against serious limitations.

The aggregation and dissemination of dispersed information has, of course, been a central theme in Austrian economics. Hayek (1945) famously pointed out the function of institutions as communication mechanisms. There is wide scope for the application of Austrian concepts in intelligence and counterintelligence. To make a start on such a research program, we focus on the emerging scholarship concerned with improving decision-making under risk and uncertainty within intelligence agencies, especially applications of behavioural decision theory. A prominent theme of this literature is improving probability estimation and communication and mitigating the type of decision-making biases identified by Tversky & Kahneman (1974). Consider Friedman & Zeckhauser's (2015, pp.77–78) account of deliberations preceding the 'kill bin Laden' operation:

Throughout Spring 2011, intelligence analysts and officials debated the chances that Osama bin Laden was living in Abbottabad, Pakistan. This question pervaded roughly 40 intelligence reviews and several meetings between President Obama and his top officials. Opinions varied widely. In a key discussion that took place in March, for instance, the president convened several advisors to lay out the uncertainty about bin Laden's location as clearly as possible....The CIA team leader assigned to the pursuit of bin Laden assessed those chances to be as high as 95 per cent, Deputy Director of Intelligence Michael Morell offered a figure of 60 per cent, and most people seemed to place their confidence level at about 80 per cent though some were as low as 40 or even 30 per cent. Other versions of the meeting agree that the president was given something like this range of probabilities, and that he struggled to interpret what this meant. President Obama reportedly complained that the discussion offered not more certainty but more confusion, and that his advisors were offering probabili-

¹ Department of Defence news briefing, February 12, 2002.

ties that disguised uncertainty as opposed to actually providing you with more useful information. Ultimately, the president concluded that the odds that bin Laden was living in Abbottabad were about 50/50.

Here, the problem is framed as one deriving from the challenges inherent in probability estimation and communication. The implication that runs through the relevant intelligence scholarship, including Friedman & Zeckhauser (2012), Dhimi et al., (2015), Barnes (2016), Dhimi (2017) and Irwin & Mandel (2019), is that problems such as this can be overcome with better probability estimation techniques and, even more importantly, better communication. In contrast, an Austrian perspective would begin by recognising the knowledge problem and the presence of true Knightian uncertainty vis-à-vis risk. Austrian analysis of entrepreneurial discovery and judgement under conditions of uncertainty can enrich intelligence scholarship while avoiding fostering a belief that if only one could tidy up certain aspects of probability assessment and communication, one could always assess the odds perfectly and could always make the best decision.²

Another classic problem in intelligence concerns the possibility of correctly anticipating the decision of an adversary. That this is believed to be possible is evident from the explanation that is offered for a failure to do so. This explanation is called *mirror-imaging*. Mirror-imaging has been offered as an explanation for the repeated failures of Western intelligence agencies to anticipate the Soviets, including the Soviet decision to invade Czechoslovakia in 1968.³ Of course, this explanation implies that it is possible, if one is careful, to correctly forecast the decisions of another party. There are severe limits to such an undertaking that are not clearly elucidated by either orthodox or behavioural decision theory. In fact, decision theory in both its forms appears to imply that it is quite possible to forecast an adversary's choice provided one has a full list of the adversary's alternatives and can adequately determine the possible outcomes and their probabilities. Austrian analysis of entrepreneurial discovery and judgement draws our attention to the substantial differences that characterise the perceptions of different individuals, even when they confront similar problems.

In this paper, we examine some contemporary developments in intelligence and counterintelligence, especially the 'big data' information problem and the use of algorithms and artificial intelligence (AI) to manage it. We explain how decision theory relates to these developments and how intelligence scholars are using orthodox⁴ and behavioural⁵ decision theory to improve decision-making. We explain how Austrian analysis can complement these efforts. True uncertainty, action in the face of it, and

² No-one would mean to imply that we could always 'find bin Laden' or solve other such problems with perfect accuracy and zero error. Expected utility maximisers do not always get the outcome they most desire. The implication is, that even if the desired outcome was not achieved, if we could tidy up our processes sufficiently, we would always or more often choose the best course of action to take *ex ante*.

³ See Baroch (1982).

⁴ The standard representatives of which are von Neumann & Morgenstern's (1947) expected utility theory and its variants (see Fishburn 1989).

⁵ Behavioural decision theory adjusts the expected utility theory of von Neumann & Morgenstern. Kahneman & Tversky's (1979) prospect theory, for example, incorporates several behavioural factors: reference points, loss aversion, and probability weighting that explain how human decisions may diverge from the prescriptions of expected utility theory.

its resolution by processes of discovery and entrepreneurial judgement have occupied the attention of the Austrian school for decades. Austrian concepts such as time, ignorance, knowledge, discovery, and entrepreneurial judgement, give us a deeper appreciation of the decision-making process than can be achieved through orthodox and behavioural decision theory alone. Austrian analysis charts a way forward for the decision-maker once it is accepted that the problem is not incomplete information to be resolved by costly search but true uncertainty.

2 Information & Algorithms: an illustrative example

Intelligence is knowledge of one's enemy.⁶ While romantic notions can be attached to intelligence and counterintelligence,⁷ there are many mundane tasks that must be performed. So mundane in fact that some definitions of intelligence might have been written about any information gathering and processing activity:

Intelligence is a corporate capability to forecast change in time to do something about it. The capability involves foresight and insight and is intended to identify impending change which may be positive, representing opportunity, or negative, representing threat (Breakspear, 2013, p.678).

Stout & Warner (2018, p.517) assessed the situation as follows:

Over the last quarter century, Anglo-American scholars have sought by fits and starts to define 'intelligence', debating what it is and what it is not. Over the past decade, a rough consensus has emerged, even if such agreement does not apply across all topics treated by intelligence scholarship. Most and perhaps all practitioners and students in the field agree that intelligence helps leaders acquire and utilize information against competitors.⁸

The enormous amount of data that intelligence agencies now confront coincides with advances in algorithms and artificial intelligence that may help manage it (Brantly, 2018; Korać & Cica, 2018; Regens, 2019; Horowitz & Kahn, 2021; Spoor & Rothman, 2021). Like its role in information systems design,⁹ decision theory can be found at the heart of these developments.

Decision theory is the economics of information. It aims to optimise systems design such that the organisation reaches the optimal decision given all possible states of the world (Marschak & Radner, 1972). From the beginning, decision theory had strong links to the defence sector, with important early work produced at

⁶ Troy (1991). Also see Warner (2002). There are five sources of intelligence: open source (OSINT); human intelligence (HUMINT); signals intelligence (SIGINT); measurements and signatures intelligence (MASINT); and geospatial intelligence (GEOINT) (Stottlemire, 2015).

⁷ On the distinction between intelligence and counterintelligence, see Barnea (2017), Dempster (2020), Ehrman (2009), and Kalaris & McCoy (1988).

⁸ The process is traditionally represented as the intelligence cycle: requirements (priorities), collection, processing (collation), analysis, dissemination, new priorities (Breakspear, 2013, p.680). See Hulnick (2006) for a critique.

⁹ On decision theory and information systems see Marschak (1971). On analytical standards in the intelligence community see Reinhold et al., (2020).

RAND.¹⁰ The ubiquity of decision theory, its reach into every nook and cranny of organisational decision-making structures, has made it a part of organisational routine (Nelson & Winter, 2002). As algorithms and AI become the latest technological development to be applied in the pursuit of more data and better decision-making, it is more important than ever to understand the place of decision theory in the design and application of these technologies because the human decision-making processes that underpin them have never been more easily obscured.

To help us understand how algorithm design can be guided by decision theory and how algorithms can subsequently transform decision-making in the contexts to which they are applied, we can use a prominent example from financial economics. This is Markowitz' (1952; 1959) critical line algorithm. The algorithm was developed in the early 1950s. It subsequently prompted and guided radical changes to portfolio management¹¹ and came to represent the cornerstone of computational finance. It is also a stark illustration of the multi-level role that decision theory plays in algorithm design.

On one level, Markowitz's algorithm illustrates the architecture of the problem space within which an algorithm works to produce its output (i.e., a set of efficient portfolios). On a deeper level, it highlights the architecture within which the designer works (i.e., the operations research framework¹² that guided Markowitz' design of his critical line algorithm). Like information systems design, the design of an algorithm to identify the optimal choice within a certain problem space¹³ is a meta-problem that can itself be analysed as a prioritisation problem with an optimal solution.

The use of algorithms and AI¹⁴ by intelligence agencies involves decisions that are only just beginning to be analysed with the help of behavioural decision theory (Phillips & Pohl, 2020). In fact, as Phillips & Pohl (2020) highlight, the human dimension of algorithm and AI design is too easily overlooked. For example, writing for the National Security Agency's (NSA) magazine, Kegelmeyer (2019, pp.12–13) states that the two vulnerabilities of machine learning (an application of AI) are: (1) similarities between the 'training' data and 'test' data that conspire to give a false impression of the accuracy of the system; and (2) inaccurate 'ground truth' labels.¹⁵ Both vulnerabilities are represented as mere technical issues when they are in fact human

¹⁰ A classic example is Hitch (1953), who illustrates his points with examples based on military missions, ship convoys, and destroyers.

¹¹ And advances in finance theory, especially asset pricing (see Sharpe (1964), Lintner (1965) and Mossin (1966)).

¹² Markowitz' work at RAND, his associations with Jacob Marschak, a pioneer in decision theory and information systems whose work we have already referenced, and George Dantzig, a pioneer in the use of optimisation techniques, including the development of the simplex algorithm, helped him to transform the portfolio problem into one that is amenable to algorithmic computation.

¹³ The algorithm and AI design will maximise the expected utility of the algorithm's outputs or AI's 'decisions' as measured by the preferences of the designers (Lineal, 1994; Huberman & Hogg, 1995; Nisan & Ronen, 2001).

¹⁴ An algorithm is a series of steps that computers can be programmed to follow to reach a solution (e.g., efficient portfolio). Artificial intelligence (AI) is a system where the computer, based on the solutions generated by the algorithms, implements decisions (e.g., trading continuously to keep a portfolio efficient).

¹⁵ For example, if a person categorises documents by hand into topics, the resulting categorisation represents that person's 'ground truth'. It is expected that the algorithm will produce a similar categorisation. If the ground truth labels are 'wrong', the algorithm might not perform well, or it might perform well on

decision problems. Intelligence scholars seeking to improve decision-making within intelligence agencies (e.g., Whitesmith 2020) are rushing to catch up with these new developments (Phillips & Pohl, 2020, 2021).

Behavioural decision theory warns us, for instance, that design choices may be shaped by the ease with which a designer can call to mind a problem that has been resolved by a particular type of algorithm (availability heuristic) or the degree to which certain features of a problem are viewed as representative of its inner structure (representative heuristic).¹⁶ An intelligence agency whose algorithm designers see a problem in a certain way may be at risk of projecting this onto their opponents' perception of the problem. Efforts to counter a rival's AI system may be undermined. This is a modern *mirror-imaging* problem. Much can be gained by consciously taking a decision theory perspective to this and the other challenges confronting the intelligence community. However, our contention is that given the fundamental similarities between orthodox and behavioural decision theory, there are limits to the gains that can be made by drawing on behavioural decision theory alone. These limits can, at the very least, be better appreciated from an Austrian perspective.

3 True uncertainty, Discovery, & Pandora's Box of Intelligence

Orthodox and behavioural decision theory have a simple underlying structure consisting of a 'problem space' populated by a set of alternatives (from which the decision-maker chooses), possible outcomes (x 's), and probabilities (p 's). The problem space is delineated in the sense that the probabilities sum to one across all outcomes. There are no surprises. Knight's (1921) distinction between risk and uncertainty, together with Ellsberg's (1961) contribution, should lead us to conclude that orthodox and behavioural decision theory will run into problems when the p 's do not sum to one because the decision-makers don't even know what the p 's are (or the x 's for that matter). However, formal work is often solely focused on risk rather than true uncertainty or assumes that uncertainty can be reduced to risk through the assignment of subjective probabilities (Savage, 1954; Anscombe & Aumann, 1963).

Dealing with uncertainty vis-à-vis risk is an emerging research program in intelligence scholarship (Friedman & Zeckhauser, 2012). That there might be a mathematical solution to the uncertainty problem is a discernible theme in the literature (see Javorsek & Schwitz 2014). This may stem from the influence of decision theory and the closely related discipline known as search theory (see Stigler 1961). Search theory deals with the way in which the problem space is populated. That is, how the decision-maker finds the alternatives, with their x 's and p 's. One explanation is Weitzman's (1978) Pandora's Box approach. Here, Weitzman (1978) depicts alternatives as distributed around the problem space in separate unopened boxes with a known probability distribution, F_i , governing the likelihood that a particular alter-

the wrong grounds. One person's categorisation decision might diverge from another's even if they are looking at the same documents.

¹⁶ Tversky & Kahneman (1974) introduce the availability and representativeness heuristics.

native with a particular expected value will be found in a particular box. Hence, the problem space is constituted of unopened boxes, each containing an alternative.¹⁷

The decision-maker in Weitzman's (1978) problem space must pay something, c_i , to open a box i . The decision-maker has a fallback alternative with a value, v , in hand. This can always be collected (Weitzman, 1978, p.5). The decision-maker decides to open the (first) or next box. An alternative, g_i^o , is revealed.¹⁸ The decision-maker can only choose one alternative. After opening each box, the decision-maker must decide whether to continue opening boxes or choose the best alternative that has been uncovered so far or take the fallback option. The decision-maker assigns a reservation value to each unopened box. The reservation value, r_i , for each unopened box should be the lowest value for which the decision-maker can do as well by taking v , as by opening box i and choosing the revealed alternative (Olszewski & Weber, 2015, p.430). A box should only be opened if:

$$r_i \geq -c_i + Emax [v, g_i^o]$$

According to Weitzman (1978), the decision-maker should first assign reservation values to each box and then order those boxes from highest reservation value to lowest. If it makes sense to open at least one box, the decision-maker proceeds to pay the price and reveal the alternative g_i^o . The decision-maker should keep opening boxes, revealing alternatives, until the above inequality is violated. Representations of the search process with analogous structures can be found throughout the literature. The approach is generalisable to any problem that can be mapped into its structure.

Uncertainty, to the extent that it exists in Weitzman's world, is reduced to risk expressed by F_i . Since F_i is likely to be subjective,¹⁹ based on the experience, intuition, or imagination of the decision-maker, the mechanical depiction of the search optimisation process obscures the human decision-maker. This leads us to the Austrian critique of neoclassical microtheory. Kirzner (1997, p.70) explains:

Austrian theory thus diverges sharply from the notion of the individual decision that constitutes the analytical building block of neoclassical microtheory. For neoclassical microtheory each decision, whether made by consumer, firm, or resource owner, [we might add algorithm, AI, or computer], is made within a definitely known

¹⁷ During the search for bin Laden, the decision-makers first had to 'find' the alternatives (geographic locations), each of which was associated with some chance, possibly zero, of finding bin Laden. We might say that the geographic locations were contained in intelligence reports (Weitzman boxes). The search begins with the set of unopened intelligence reports. The reports (boxes) can be assigned a value based on their source etc. Once opened, a report reveals a location with a range of x 's (including, or not, finding bin Laden). Which report should be opened? And once revealed, should the alternative be chosen, or should the decision-maker move on to the next report?

¹⁸ Each g has an expected series of x 's. The x 's are nested within the g 's. The superscript denotes 'opened' (revealed).

¹⁹ For example, assume that the 'boxes' are books that an academic believes contain alternative potential research opportunities. The problem is to find the 'best' alternative by opening boxes (reading books). The best alternative is in one of the books that lay before the academic. The problem space is delineated and there are no 'other' books. The probability distribution is based on the academic's experience. Certain publishers, authors, time periods etc. may be more or less likely to represent the 'box' containing the best alternative.

framework made up of a given objective function, a given set of resource constraints, and a given set of technologically or economically feasible ways of transforming resources into desired objectives. Uncertainty, while of course recognized as surrounding each decision, expresses itself in the form of known probability distributions relating to the given elements of this known framework.


In place of this mechanical search and decision process taking place within a delineated problem space, the Austrian school has traditionally recognised the non-delineated nature of most important real-world problems and explores human action in the face of this uncertainty. Kirzner (1997, p.62) makes a point that is important given the positive intention of the present paper:

What stamps the entrepreneurial discovery approach as Austrian is not these criticisms themselves, but rather the specific positive elements of the approach. These positive elements focus on the role of knowledge and discovery.

The knowledge Kirzner refers to is ‘un-thought-of knowledge’ or ‘unknown ignorance’. Discovery is the process by which unknown ignorance, which might be described as ignorance of un-thought-of knowledge or even as Rumsfeld’s unknown unknowns, is resolved.²⁰ This is something quite distinct from imperfect information or ‘known-to-exist information’ (about alternatives, x ’s, and p ’s), which is resolved in orthodox decision theory by costly search (see Kirzner 1997, p.65). The distinction between “knowledge and discovery” vis-à-vis “imperfect information and costly search” leads to the introduction of the central character in Austrian economics, the entrepreneur.

The decision-maker from orthodox and behavioural decision theory traverses a problem space that embeds all alternatives and, therefore, all possible outcomes that can be experienced. The probabilities assigned to each possible outcome sum to one. This is illustrated by the left-hand panel of Fig. 1. Here, x ’s are uncovered by costly search, but the problem space is delineated. It is not expanded by the costly

Fig. 1 A Delineated Problem Space (Left) Leaves No Scope for Discovery



²⁰ Also see Hayek (1968; 2002).

search process. The Austrian perspective championed by Kirzner (1973; 1985; 1997) stands in stark contrast. Rather than uncovering x 's by costly search, the problem space itself is expanded by a discovery process. This is illustrated by the right-hand panel of Fig. 1. The problem space is open or non-delineated and what lays beyond the decision-maker's current knowledge is shrouded by true uncertainty. Where the delineated problem space may be filled in by simply deploying enough manpower or computing power, the non-delineated problem space cannot be approached in such a purely mechanical way.

The distinction between these two types of problem spaces takes us a step closer to a deeper understanding of the structure of practical problems encountered by intelligence agencies. Clyde Conrad was a non-commissioned officer (NCO) operating in the US Army's 8th Infantry Division in Bad Kreuznach, West Germany. He was also a mole. Sometime in the late 1970s, CIA officers began receiving reports from their sources behind the Iron Curtain that the Hungarians had an agent somewhere inside the US Army. The Army's Foreign Counterintelligence Activity (FCA) began an investigation in the 1980s. They started by looking for an NCO. Given the duration of the espionage activity, a commissioned officer, with short-term rotations through West Germany, would be unlikely to be involved. The NCO must also have had continuous access to the right classified material (i.e., the documents that were turning up in Hungary) and might have a Hungarian family connection or heritage. Finding the right person seems straightforward and only a matter of time but as Herrington (1999, pp.94–95) explains:

Since the 1970s, some twenty-five to thirty thousand soldiers—thousands of whom were NCOs—had served in Germany in duty positions with access to such documents. Determining who these people were, where they might be in 1985, and whether or not they met other elements of the profile was formidable in the days before the proliferation of automated systems. Even a task as seemingly simple as determining possible Hungarian ethnic background was easier said than done. If a soldier was born in Hungary, that fact would appear at a certain place in his records. But if the soldier was not born in Hungary but his father was Hungarian, investigators were reduced to the crude and unreliable method of searching for Hungarian-looking names. Worse yet, if a soldier's mother was Hungarian and his father's name was, say, O'Brien, FCA investigators had no easy way of detecting the Hungarian background on his mother's side.... Month after tedious month, the persistent FCA agents worked long hours exploring every possible avenue that might lead to their man.

It is true that the final phase of the problem consisting of searching through the records of 35,000 soldiers for red flags can be structured to some extent as a Weitzman box problem. And behavioural decision theory might shed light on specific investigative shortcomings, such as the decision biases that might explain the delay in commencing the investigation. But at the beginning of the process, many years before the identification of Conrad, there were only vague reports of the existence of a mole. And before then, of course, nothing. There was a transition from this situation to confirmation of a mole's existence to the identification of Conrad as that mole. Rather than the mechanics of the search, thinking about the transition as a discovery process deepens our perspective and leads us to consider the nature of the judgements that were brought to bear at each step, from the initial realisation of the

significance of source reports, to their meaning, to their reliability and veracity, to the focus on NCOs, to operational considerations that led to the arrest of Clyde Conrad. The uncertainty about the existence of a mole and his or her identity did not resolve itself. It was driven by human action. Austrian analysis, almost uniquely, has turned its attention to uncovering the nature of uncertainty-resolving human action.


4 Uncertainty resolution in Intelligence Agencies

While ‘discovery’ leads us to a new appreciation of the problem space, efforts to improve intelligence decision-making are drawing on quite fundamental concepts from decision theory. What is required at this early stage, therefore, is a steppingstone that connects those fundamental concepts and Austrian analysis. A likely candidate is found in the Foss-Klein framework. While emphasising entrepreneurial *judgement* over entrepreneurial *discovery* may not be to the liking of all Austrian economists (see Sautet 2018), the concept of judgement is more familiar to scholars who are using orthodox decision theory to think about problems in intelligence. As such, for the context that we have been exploring, Foss & Klein’s (2012) and Packard et al., (2017) emphasis on the role of judgement in the resolution of uncertainty provides a more direct connection between decision theory and Austrian analysis of the type of problems confronting intelligence agencies.²¹ It represents a recognisable complement to orthodox analysis and, in some ways, prepares the ground for a more nuanced application of orthodox and behavioural decision theory.

In orthodox decision theory, uncertainty is reduced to risk by identifying a set of states of the world, say 1, 2, 3, 4 and 5, one of which the decision-maker will eventually experience. This set of states of the world is ‘partitioned’. The elements of each partition are events. At first, the decision-maker has no information. The partition confronting the decision-maker is simply one big collection of all possible states of the world. As time passes, the decision-maker observes events and watches as the partition of events becomes finer and finer until each partition contains one state of the world and one of those states is realised. This is illustrated by Fig. 2.

Fig. 2 The Resolution of Uncertainty in Orthodox Decision Theory²²

²² See Lengwiler (2004, p.12)



²¹ We thank the anonymous reviewer for the guidance in this regard.

Such a representation of uncertainty may be logically permissible in situations where the decision-maker cannot influence the course of events. For most interesting decision problems, though, the passive resolution of uncertainty won't do. We explained with reference to the Clyde Conrad case (and there are many similar examples and many analogous problems), that the inherent uncertainty of the context would not have resolved itself. Indeed, human decision-makers need something to be uncertain about before that uncertainty can be resolved. The recognition of the possibility of an intelligence mole precedes everything else in this example. If we want to explain how uncertainty is actively resolved by human decision-makers and, in the process, extend the scope of ongoing efforts by intelligence scholars to enhance decision-making, we require a deeper explanation that places an active decision-maker at the heart of the matter. Foss & Klein (2012) and Packard et al., (2017) explain how the entrepreneur exercising judgement (actively) resolves total uncertainty.

When characterised by total uncertainty, the problem space is unpopulated by alternatives, outcomes, and probabilities. Under these most difficult of circumstances, orthodox and behavioural decision theory do not have much to say about how the decision-maker proceeds. Packard et al., (2017) explain how the decision-maker reels the problem in, rendering it more akin to the types of problem spaces that orthodox decision theorists are used to dealing with (also see Baron 2004; 2006 and Klein 2008). Packard et al., (2017, p.10) identify three types of uncertainty confronting the decision-maker: (1) total uncertainty, where nothing is delineated and the problem space is unpopulated; (2) environmental uncertainty, where the set of alternatives is delineated but the set of outcomes (x 's) is non-delineated (see Fig. 1, righthand panel); and (3) creative uncertainty, where the set of outcomes (x 's) is delineated but the set of alternatives is non-delineated.

The decision-maker facing total uncertainty, takes steps to resolve it first into creative uncertainty by addressing environmental uncertainty or vice versa. If environmental uncertainty is addressed first, the set of outcomes is delineated before the set of alternatives is delineated. If creative uncertainty is addressed first, the set of alternatives is delineated before the set of outcomes is delineated. Packard et al., (2017, p.18) suggest that different types of reasoning are associated with each approach (also see Sarasvathy 2001). Addressing environmental uncertainty first (delineating the outcomes) is associated with causal reasoning. Addressing creative uncertainty first (delineating the alternatives) is associated with effectual reasoning.

Causal reasoning trails problem recognition. Packard et al., (2017, pp.18–19) argue that the decision-maker, upon recognising a problem, becomes concerned with judging the potential outcomes, x 's, of *the problem's solution* in terms of *value* and *viability*. For example, the intelligence agency decision-maker may have recognised that the background identities manufactured for agents are now much easier for rival agencies to detect using sophisticated technology (Lucas, 2019). How valuable would it be to solve this problem? Notice that this question, whether answered in the affirmative or the negative, effectively closes the problem space by delineating the set of x 's as 'valuable' or 'not valuable'. The set of alternatives, however, remains open. Creative uncertainty remains to be resolved. The decision-maker now works on finding a viable alternative given the available resources, eventually hitting upon a solution to the problem, or discovering that there is, for now, no feasible solution.

Unlike the decision-maker from orthodox decision theory, the decision-maker does not order a pre-existing set of alternatives and, having decided to choose one, take as the risk a range of possible x 's. Rather, the x 's are actively delineated, resolving the environmental uncertainty. Then, the alternatives are actively delineated, resolving the creative uncertainty. The result of this process is either an alternative to pursue or the confirmation that there is no feasible alternative.

Effectual reasoning is separated from problem recognition. The decision-maker following an effectual reasoning approach transitions out of total uncertainty by first addressing creative uncertainty (Packard et al., 2017, p.19). Rather than problem recognition, the decision-maker starts with a recognition of the value of available resources. The decision-maker therefore focuses on delineating the set of alternatives, resolving the creative uncertainty first. Once the alternatives have been delineated, the decision-maker turns his or her attention to delineating the set of outcomes and the resolution of the environmental uncertainty. The result of this transition from total uncertainty to a delineated problem space is either a valuable and viable alternative or the process begins afresh. Both the causal and effectual approaches gradually delineate the alternatives and outcomes. The decision-maker pursuing either of these approaches actively resolves uncertainty. It does not resolve itself.

Packard et al., (2017, p.28) make an interesting observation towards the end of their paper. They suggest that whether the decision-maker takes the causal or effectual approach depends on how they perceive the problem space. If the decision-maker perceives the set of outcomes to be delineate-able, causal reasoning may be favoured. If the decision-maker perceives the set of alternatives to be delineate-able, effectual reasoning may be preferred. This raises a question that does not emerge from orthodox and behavioural decision theory. That is, why does someone see the problem space in a particular way or, more fundamentally, how do people form perceptions of problem spaces? Packard et al., (2017) suggest that inexperience associated with a poor understanding of uncertainty may produce a counterproductive degree of overconfidence in their ability to delineate the problem space. This represents an entry point for the type of behavioural decision theory that intelligence scholars seek to use to improve decision-making.

5 Information, Imagination & Surprise

Intelligence agencies can be surprised. China has been focused on AI for a decade (Stavridis, 2021) and Vladimir Putin notably stated in 2017 that whichever country becomes the leader in AI will become 'ruler of the world' (Petrella et al., 2021). But a stunning report released in 2021 by the National Security Commission concluded that, "America is not prepared to defend or compete in the AI era" (National Security Commission 2021, p.1). Then, in October 2021, China tested a hypersonic missile that reportedly caught US intelligence agencies completely off-guard (Marcus, 2021). This was compounded in early 2022 when Russia claimed to have tested a similar capability with its long-range hypersonic Zircon cruise missile (Dutton, 2022).

These examples indicate in the most serious possible way that a problem space can, at any moment, be transformed completely by the revelation of an adversary's

new capability or the deficiency of one's own relative capability. It is also a potent reminder of the types of information signals that shape the decisions of non-price institutions like departments of defence. In the case of hypersonic missiles, the United States was obviously not unaware of the possibility that such technology could exist but seems to have been unaware that such capability was well within reach of its adversaries and certainly unaware that the technology could so soon play a role, even for posturing purposes, in an actual theatre of war. Only the demonstration by an adversary provided the necessary information. Should the United States pursue hypersonic missile technology? If so, at what cost? Once developed, how many missiles should be produced? Twice as many as the Russians and Chinese combined? Four times as many? It comes down to bargaining between and within departments and agencies, with political interests pushing things now one way, now the other.

Intelligence agencies cannot rationally calculate how much intelligence to produce, and its consumers cannot rationally calculate how much intelligence to demand. Compare this to a market setting, such as the financial markets. Here, information is anything that is relevant to the market value of tradeable securities. Whatever has no such value, though it may appear to have the characteristics of information, is called noise. Both noise and information are reflected in market prices. Over time, information traders are rewarded, and noise traders are punished. Prices transmit this information to every market observer, allowing them, in turn, to make rational investment decisions. Market participants have worked out what information is most relevant and concentrate their efforts on gathering it up to the point at which its marginal cost and marginal benefit are equal. While information is the stock in trade of intelligence agencies, there is no equivalent price mechanism to guide its collection. Nor is there an equivalent price mechanism guiding those who consume it or act upon it.

And information gathered and reported by intelligence agencies is only a part of the flow of information that constantly ripples through the defence sector. Information comes from other sources, including demonstrations of capability by adversaries. Without price signals guiding the volume and direction of flow, decision-makers rely on other structures (organisational, hierarchical, political, etc.). It is no surprise, given that technological developments can be driven centrally, that technological innovation has become something of a hallmark of the modern intelligence services. At the end of the day, however, chaotically produced torrents of "bits", whether information, disinformation, misinformation, or noise, is in no-one's interest and the intelligence community is aware of the need to improve the rationality of decision-making, even though they may not be aware of the inherent irrationality that must characterise a non-market context. Can orthodox and behavioural decision theory together with Austrian analysis (and any other useful theoretical frameworks at our disposal) help decision-makers prepare for and navigate situations in such a context and, specifically, in cases where the problem space rapidly changes?

The technological developments that are transforming many aspects of human activity, not just intelligence collection and analysis, embed multiple layers of human decision-making. The choices that designers make are usually guided by attempts to delineate a problem space such that an algorithm and the AI system that may be embedded within it can function. Our classic example was Markowitz' critical line algorithm, which turns all investments and portfolios into 'mean-variance pairs' to

facilitate a quadratic programming optimisation solution to the ‘portfolio problem’. The study of these choices can be deepened by considerations of entrepreneurial judgement. Algorithm and AI design is not a purely mechanical process. Once a problem is recognised, true uncertainty may engulf the decision-makers in their efforts to find a technical solution. Or, conversely, given the power of newly available software and hardware, and the volume of ‘big data’, decision-makers may be completely uncertain as to the best way to deploy it.

One of the most interesting aspects of algorithm and AI systems is their potential to surprise even those who develop them. Google, for instance, developed an algorithm that would steer an unmanned helium balloon from Puerto Rico to Peru. During its flight, the balloon repeatedly veered off-course and had to be manually put back on track. However, the developers later discovered that the AI system had ‘learned’ to “tack” or zigzag into and out of the wind (Baraniuk 2021). The system had ‘learned’ this autonomously. It was not programmed explicitly into the algorithm as a method by which the AI could traverse the geographical problem space. Another example is He et al.’s (2019) ‘AI Universe Simulation’, developed to produce simulations of the universe. Surprisingly, the AI could produce simulations of the universe under conditions that it had not been trained to consider. One of the authors explained the surprise, “It’s like teaching image recognition software with lots of pictures of cats and dogs, but then it’s able to recognize elephants” (Science Daily 2019). A surprising outcome is distinct from the uncertainty from which it may emerge. Nevertheless, the possibility of surprising outcomes can be encompassed within the type of frameworks that we have been discussing throughout this paper.

G.L.S. Shackle’s (1961; 1969) work has long attracted the interest of Austrian economists (e.g., Lachmann 1976). Packard et al., (2017) refer to it extensively. While Shackle also provides an explanation of how the decision-maker proceeds in the face of true uncertainty, it is his non-distributional measure of uncertainty, called ‘potential surprise’, that catches our attention here. Imagine a decision-maker attempting to determine what *will* happen in any of the quickly evolving, technologically sophisticated contexts that we have described. Orthodox decision theory suggests that the various outcomes, x ’s, that the decision-maker determines to be possible are assigned (subjectively) probabilities that sum to one. A failure to foresee one of the x ’s is a mistake that will not be systematically repeated and, furthermore, although the decision-maker may not have anticipated an outcome, the outcome was still embedded within the problem space. The decision-maker just did not find it. Shackle recognised the limitations of a probability measure that must sum to one.

Let x be the event “rain”, then not- x is “not rain”. We can divide “not-rain” into all the other weather events (hail, fog, sun...). Shackle (1961; 1969, pp.73–74) asks:


How can probability deal with this need to assign equal status to “rain” and “not rain”? To give a probability of 0.50 to rain will leave 0.50 to be shared amongst sun, fog, snow and hail, and if, as may well be the case, he feels that each of them deserves an equal status with rain, he will feel totally frustrated. By contrast, potential surprise, because it is non-distributional, can be assigned in zero degree to an unlimited number of particularised components of “not rain” and so to “not rain” itself.

Rather than determine what *will* happen, Shackle’s decision-maker determines what *can* happen. He or she can say how surprised they would be if something were to happen. Shackle develops a non-distributional measure of uncertainty based on the degree of potential surprise that the decision-maker attaches to *imagined* outcomes. An outcome that is perfectly possible has a zero degree of potential surprise while an outcome that the decision-maker believes to be impossible is accorded the maximum degree of potential surprise. All other outcomes are given a degree of potential surprise in the range between zero and the maximum. The only constraint placed on potential surprise as a non-distributional measure of uncertainty is that when it is applied to an outcome and its contradictory, one or other of these must be accorded a zero degree of surprise. Since many outcomes can be accorded the same degree of potential surprise (zero, the maximum or anywhere in between), there is no summing to one. Potential surprise is, in Shackle’s words, like a net that encompasses imagined outcomes (x ’s). If the net encompasses what turns out to be the actual experienced outcome, the decision-maker’s judgement was correct.

At a basic level, an algorithm or AI system that performs unsurprisingly can be said to have either ‘worked’ or ‘failed’. It either traverses the delineated problem space as expected, solving whatever problem it was designed to solve, or it fails to do so, either at all or as effectively as the designers hoped. Whether it works or fails, its behaviour can be described within the delineated problem space of orthodox and behavioural decision theory. The possibility that an algorithm or AI system can (not will) surprise either its creators or the adversary who learns of it, Shackle’s unique depiction of an unfolding problem space, open to both imagination and surprise, may be used to represent the situation (Fig. 3).

One of Shackle’s insights regarding this type of non-delineated unfolding problem space is that while it is open, it is not arbitrary. In recent years, theoretical physicists have begun to draw on decision theory to refine their analysis of ‘quantum foundations’. They face the analogous problem of modelling an open but non-arbitrary future (see Smolin 2019, pp.198–204). Shackle’s open problem space includes not just ‘sum to 1’ probable x ’s but *possible* x ’s. As Smolin explains (p.199), Werner Heisenberg recognised that possibilities must be included in ‘reality’ because, “...the physics of the possible influences the future of the actual.” As such, an interpretation

Fig. 3 A Delineated Problem Space vis-à-vis Shackle’s Unfolding Space



of quantum mechanics based on possibilities also implies the future is open. Open, but not arbitrary.

Just as Shackle argues that the world is not arbitrary, that there are boundaries around the possible, Smolin (p.199) refers to the idea of ‘adjacent possible’. Possible near-future events are not yet real but define what might be real. In the classic ‘Schrödinger Cat’ thought experiment, the adjacent possible includes ‘live cat’ and ‘dead cat’. The adjacent possible, as Smolin goes on to say, does not include ‘brontosaurus’ or ‘dog’. In contexts where problem spaces are open or can be opened, perhaps explosively, by surprising innovations, the decision-maker is not completely adrift. To understand the context and prepare accordingly within it requires different perspectives than those that are found in orthodox theory. But what does ‘prepare accordingly’ in the face of true uncertainty and potential surprise mean?

We can look for guidance in contexts that may portend our own future. In finance, technology continuously reshapes both theory and practice. Currently, developments are moving so fast that academic research is struggling to keep up (Huck, 2019). In an insightful paper written for practitioners, O’Hara (2014) argues that big data, algorithms, and AI have necessitated a deepening of analysis and decision-making in finance, moving beyond data gathering to strategic consideration of the reach of the new technologies. In one of the many potential overlaps between high-end finance and intelligence, O’Hara (2014, p.23) points out the importance of ‘being invisible’ to the algorithms and AI. Human decision-makers cannot gather or process information as quickly as a computer, but they can ensure that the ‘weaknesses’ that the algorithms and AI are programmed to exploit are eliminated. This requires an understanding of the problem space within which the algorithm operates, the intentions of the designer, and recognition that the new technology might have completely surprised the human who designed it. Such a surprise lays beyond the problem space of orthodox and behavioural decision theory, leading us to suggest that the future strategic value of intelligence analysis is located beyond information acquisition, however fast and however vast. It may very well reside in the deep appreciation of the various problem spaces facing or imagined by adversaries and the technologies they have deployed within them.

6 Conclusions

Members of the intelligence community use market terminology to describe intelligence related activities. They speak of the intelligence product and of intelligence consumers. There are, however, no market-generated price signals to guide the collection, dissemination, or consumption of intelligence. Price signals are replaced by bargaining within organisational-hierarchical-political frameworks shaped by centrally directed innovation. Often, information arriving from outside of these frameworks plays a more decisive role than the intelligence gathered within them. This includes demonstrations of new capability by adversaries, such as the development of hypersonic missiles or, in the old days, Sputnik. If we were to allow ourselves to be misled by the market-derived terminology and analyse intelligence agencies as

market-based institutions instead of non-price institutions, our analysis would likely produce seriously mistaken conclusions.

The intelligence community is aware of the calculation problem faced by decision-makers in this context, though they are probably unaware that this difficulty stems in part from the absence of a genuine market for intelligence, despite the vernacular. Bringing order and rationality to the process is desirable but difficult to achieve. It comes as no surprise, given that technological innovation is something that can be centrally driven, that technological advances, in surveillance among other things, have become a defining characteristic of modern intelligence agencies. However, the performance of intelligence agencies is, at the end of the day, not judged on the volume of information collected. Whether that volume be a little or a tremendous, tremendous amount, its relevance and its rational utilisation is what matters. Recognising this, intelligence scholars have been engaged in an ongoing and expanding program that investigates the decision-making and communication processes at work in the context that we have been describing. In so doing, they have ventured into territory of the economics of information.

Under conditions of true uncertainty, orthodox and behavioural decision theory offer little guidance to the decision-maker. By contrast, Austrian analysis of discovery processes and entrepreneurial judgement focuses directly on this problem. This work comes at a time when intelligence practice is being reshaped by new information flows and new technologies that are designed to manage it. From a traditional decision theory perspective, vast new information flows merely exacerbate the imperfect information problem while technology represents new ways to carry out the costly search needed to resolve it. But the challenge confronting intelligence agencies is not one that can be overcome by 'faster and vaster' information acquisition and analysis. Rather, strategic advantage resides beyond information acquisition in a deeper understanding of the ways, including technology-based, in which rivals have set to work to delineate and navigate the problem spaces that confront them. Without price signals, though, information about missteps or advances experienced by centrally driven technological innovation does not flow freely, even within the agencies themselves. Whether this can be turned to one's advantage while mitigating the negative implications for one's own agency is a problem that, like the other aspects of intelligence and counterintelligence that we have touched on, offers much for the Austrian economist to ponder.

Funding declaration No funding was received for this paper.

Open Access funding enabled and organized by CAUL and its Member Institutions

Declarations

Conflict of interest The authors declare no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this

article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anscombe, F. J., & Aumann, R. J. (1963). A Definition of Subjective Probability. *Annals of the Mathematical Statistics*, 34, 199–205
- Baraniuk, C. How Google's Balloons Surprised Their Creator., & News, B. B. C. (2021). February 24. <https://www.bbc.com/future/article/20210222-how-googles-hot-air-balloon-surprised-its-creators>
- Barnea, A. (2017). Counterintelligence: Stepson of the Intelligence Discipline. *Israel Affairs*, 23, 715–726
- Barnes, A. (2016). Making Intelligence Analysis More Intelligent: Using Numeric Probabilities. *Intelligence and National Security*, 31, 327–344
- Baroch, C. (1982). *The Mirror-Image Fallacy: Understanding the Soviet Union*. June 29. Washington, D.C.: The Heritage Foundation
- Baron, R. A. (2004). Opportunity Recognition: A Cognitive Perspective. *Academy of Management Proceedings* 2004, 1, A1-A6
- Baron, R. A. (2006). Opportunity Recognition as Pattern Recognition: How Entrepreneurs 'Connect the Dots' to Identify New Business Opportunities. *Academy of Management Perspectives*, 20, 104–119
- Brantly, A. F. (2018). When Everything Becomes Intelligence: Machine Learning and the Connected World. *Intelligence and National Security*, 33, 562–573
- Breakpear, A. (2013). A New Definition of Intelligence. *Intelligence and National Security*, 28, 678–693
- Dempster, G. M. (2020). Why (a Theory of) Opportunity Matters: Refining the Austrian View of Entrepreneurial Discovery. *Quarterly Journal of Austrian Economics*, 23, 427–461
- Dhami, M. K. (2017). Towards an Evidence-Based Approach to Communicating Uncertainty in Intelligence Analysis. *Intelligence and National Security*, 33, 257–272
- Dhami, M. K., Mandel, D. R., Mellers, B. A., & Tetlock, P. E. (2015). Improving Intelligence Analysis with Decision Science. *Perspectives on Psychological Science*, 10, 753–757
- Dutton, J. (2022). *How Russia's Zircon Hypersonic Missiles Could Turn the Tide of the Ukraine War* (p. 30). Newsweek
- Ehrman, J. (2009). What are We Talking About When We Talk About Counterintelligence? *Studies in Intelligence*, 53, 5–20
- Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75, 643–669
- Fishburn, P. C. (1989). Retrospective on the Utility Theory of von Neumann and Morgenstern. *Journal of Risk and Uncertainty*, 2, 127–158
- Foss, N., & Klein, P. (2012). *Organising Entrepreneurial Judgement: A New Approach to the Firm*. Cambridge, UK: Cambridge University Press
- Friedman, J. A., & Zeckhauser, R. (2012). Assessing Uncertainty in Intelligence. *Intelligence and National Security*, 27, 824–847
- Friedman, J. A., & Zeckhauser, R. (2015). Handling and Mishandling Estimative Probability: Likelihood, Confidence, and the Search for Bin Laden. *Intelligence and National Security*, 30, 77–99
- Hayek, F. A. (1945). The Use of Knowledge in Society. *American Economic Review*, 35, 519–530
- Hayek, F. A. (1968). ; 2002. Competition as a Discovery Procedure. *Quarterly Journal of Austrian Economics*, 5, 9–23 [M.S. Snow, Translator]
- He, S., Li, Y., Feng, Y., Ho, S., Ravanbakhsh, S., Chen, W., & Póczos, B. (2019). Learning to Predict the Cosmological Structure Formation. *Proceedings of the National Academy of Sciences*, 116, 13825–13832
- Herrington, S. A. (1999). *Traitors Among Us*. Novato, CA: Presidio Press
- Hitch, C. (1953). Sub-Optimisation in Operations Problems. *Journal of the Operations Research Society of America*, 1, 87–99
- Horowitz, M. C., & Kahn, L. (2021). How Joe Biden Can Use Confidence-Building Measures for Military Uses of AI. *Bulletin of Atomic Scientists*, 77, 33–35

- Huberman, B. A., & Hogg, T. (1995). Distributed Computation as an Economic System. *Journal of Economic Perspectives*, 141–152
- Huck, N. (2019). Large Data Sets and Machine Learning: Applications to Statistical Arbitrage. *European Journal of Operational Research*, 278, 330–342
- Hulnick, A. (2006). What's Wrong with the Intelligence Cycle? *Intelligence and National Security*, 21, 959–979
- Irwin, D., & Mandel, D. R. (2019). Improving Information Evaluation for Intelligence Production. *Intelligence and National Security*, 34, 503–525
- Javorsek, I. I., D., & Schwitz, J. G. (2014). Probing Uncertainty: Complexity, and Human Agency in Intelligence. *Intelligence and National Security*, 29, 639–653
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47, 263–291
- Kalaris, G., & McCoy, L. (1988). Counterintelligence for the 1990s. *International Journal of Intelligence and Counterintelligence*, 2, 179–187
- Kegelmeyer, W. P. (2019). Adversarial Issues in Machine Learning. *The Next Wave*, 22, 11–14
- Kirzner, I. M. (1973). *Competition and Entrepreneurship*. Chicago, IL: University of Chicago Press
- Kirzner, I. M. (1985). *Discovery and the Capitalist Process*. Chicago, IL: University of Chicago Press
- Kirzner, I. M. (1997). Entrepreneurial Discovery and the Competitive Market Process: An Austrian Approach. *Journal of Economic Literature*, 35, 60–85
- Klein, P. G. (2008). Opportunity Discovery, Entrepreneurial Action and Economic Organisation. *Strategic Entrepreneurial Journal*, 2, 175–190
- Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Boston, MA: Houghton Mifflin
- Korać, D., & Cica, D. (2018). A Mathematical Model for the Evaluation of Intelligence Products Value. *Journal of Information and Optimisation Sciences*, 39, 903–926
- Lachmann, L. M. (1976). From Mises to Shackle: An Essay of Austrian Economics and the Kaleidic Society. *Journal of Economic Literature*, 14, 54–62
- Lengwiler, Y. (2004). *Microfoundations of Financial Economics: An Introduction to General Equilibrium Asset Pricing*. Princeton, NJ: Princeton University Press
- Lineal, N. (1994). Game Theoretic Aspects of Computing. *Handbook of Game Theory* (2 vol., pp. 1339–1395). Elsevier Science Publishers B.V.
- Lintner, J. (1965). The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *Review of Economics and Statistics*, 47, 13–37
- Lucas, E. (2019). The Spycraft Revolution. *Foreign Policy*, April 27. <https://foreignpolicy.com/2019/04/27/the-spycraft-revolution-espionage-technology/>
- Marcus, J. (2021). China's Hypersonic Missile. Could it Spark a New Arms Race? *BBC News*, October 24. <https://www.bbc.com/news/world-asia-59001850>
- Markowitz, H. M. (1952). Portfolio Selection. *Journal of Finance*, 7, 77–91
- Markowitz, H. M. (1959). *Portfolio Selection: Efficient Diversification of Investments*. New York: John Wiley & Sons
- Marschak, J. (1971). Economics of Information Systems. *Journal of the American Statistical Association*, 66, 192–219
- Marschak, J., & Radner, R. (1972). *Economic Theory of Teams*. New Haven: Yale University Press
- Mossin, J. (1966). Equilibrium in a Capital Asset Market. *Econometrica*, 34, 768–783
- National Security Commission 2021. Final Report: National Security Commission on Artificial Intelligence. Arlington, VA
- Nelson, R. R., & Winter, S. G. (2002). Evolutionary Theorising in Economics. *Journal of Economic Perspectives*, 16, 23–46
- Nisan, N., & Ronen, A. (2001). Algorithmic Mechanism Design. *Games and Economic Behaviour*, 35, 166–196
- O'Hara, M. (2014). High-Frequency Trading and Its Impact on Markets. *Financial Analysts Journal*, 70, 18–27
- Olszewski, W., & Weber, R. (2015). A More General Pandora Rule? *Journal of Economic Theory*, 160, 429–437
- Packard, M. D., Clark, B. B., & Klein, P. G. (2017). *Uncertainty Types and Transitions in the Entrepreneurial Process* (p. 17). University of Omaha Marketing and Management Faculty Publications
- Petrella, S., Miller, C., & Cooper, B. (2021). Russia's Artificial Intelligence Strategy: The Role of State-Owned Firms. *Orbis*, 65, 75–100
- Phillips, P. J., & Pohl, G. (2020). Countering Intelligence Algorithms. *The RUSI Journal*, 165, 22–32

- Phillips, P. J., & Pohl, G. (2021). Tinker, Thaler, Soldier, Spy: Behavioral Economics of HUMINT Transactions and Source Prioritizations. *International Journal of Intelligence and CounterIntelligence*, 34, 17–44
- Regens, J. L. (2019). Augmenting Human Cognition to Enhance Strategic, Operational, and Tactical Intelligence. *Intelligence and National Security*, 34, 673–687
- Reinhold, D., Russo, C. M., & Eisenfeld, B. (2020). Analytical Standards in the Intelligence Community. *Journal of Strategic Security*, 14, 106–121
- Science Daily (2021). The First AI Universe Sim is Fast and Accurate and its Creators Don't Know How it Works. June 26. <https://www.sciencedaily.com/releases/2019/06/190626133800.htm>
- Sarasvathy, S. D. (2001). Causation and Effectuation: Toward a Theoretical Shift From Economic Inevitability to Entrepreneurial Contingency. *Academy of Management Review*, 26, 243–263
- Sautet, F. (2018). The Battle for the Essence of Entrepreneurship. *Review of Austrian Economics*, 31, 123–139
- Savage, L. J. (1954). *The Foundations of Statistics*. New York: John Wiley & Sons
- Shackle, G. L. S. (1961). ; 1969. *Decision, Order and Time in Human Affairs*, Cambridge University Press, Cambridge
- Sharpe, W. F. (1964). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *Journal of Finance*, 19, 425–442
- Smolin, L. (2019). *Einstein's Unfinished Revolution*. London, UK: Penguin Press
- Spoor, B., & Rothman, M. (2021). On the Critical Utility of Complexity Theory in Intelligence Studies. *Intelligence and National Security*, 36, 555–568
- Stavridis, J. (2021). Artificial Intelligence is America's Achilles Heel Against China. Bloomberg, May 21. <https://www.bloomberg.com/opinion/articles/2021-05-20/china-s-artificial-intelligence-advantage-is-america-s-achilles-heel>
- Stigler, G. J. (1961). The Economics of Information. *Journal of Political Economy*, 69, 213–225
- Stottlemire, S. A. (2015). HUMINT, OSINT, or Something New? Defining Crowdsourced Intelligence. *International Journal of Intelligence and CounterIntelligence*, 28, 578–589
- Stout, M., & Warner, M. (2018). Intelligence is as Intelligence Does. *Intelligence and National Security*, 33, 517–526
- Troy, T. F. (1991). The “Correct” Definition of Intelligence. *International Journal of Intelligence and CounterIntelligence*, 5, 433–454
- Tversky, A., & Kahneman, D. (1974). Judgement under Uncertainty: Heuristics and Biases. *Science*, 185, 1124–1131
- von Neumann, J., & Morgenstern, O. (1947). *Theory of Games and Economic Behaviour*. Princeton, NJ: Princeton University Press
- Warner, M. (2002). Wanted: A Definition of Intelligence. *Studies in Intelligence*, 46, 15–22
- Weitzman, M. (1978). Optimal Search for the Best Alternative. MIT Energy Laboratory Report MIT-EL-78-008, Department of Energy, United States of America
- Whitesmith, M. (2020). Experimental Research in Reducing the Risk of Cognitive Bias in Intelligence Analysis. *International Journal of Intelligence and CounterIntelligence*, 33, 380–405

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.