# WHITE PAPER

# DERISKING FAMILY OFFICE CAPITAL BY MAKING COMPLEXITY SIMPLE TO SEE, WITH THE BARBELL STRATEGY AND SUPPORTING HEURISTICS

Simple rules when capital gets complex

START PICNIC venture studio

startup picnic academy

16 SEPTEMBER 2026 | 16 September 2026

www.startuppicnic.co.za

By Tiisetso Maloma, Startup Picnic Venture Studio, and Startup Picnic Academy

# WHITE PAPER

## DERISKING FAMILY OFFICE CAPITAL BY MAKING COMPLEXITY SIMPLE TO SEE, WITH THE BARBELL STRATEGY AND SUPPORTING HEURISTICS

### Simple rules when capital gets complex

**16 September 2026 | www.startuppicnic.co.za**

*By Tiisetso Maloma, Startup Picnic Venture Studio, and Startup Picnic Academy*

***This guide offers a perspective for family offices. It shows why the future of capital is hard and uncertain to calculate. It offers a simple understanding of complexity and presents the barbell strategy as an overarching heuristic, with supporting heuristics for safety and upside exploration — the safe side as a refuge. It is not investment advice. It does not argue for or against any particular investment strategy. It presents a way of thinking about capital allocation under complexity and uncertainty—and at least in a simpler way.***

# CONTENTS

Contents ..... 4
Introduction: what this guide offers ..... 6
The authors ..... 7
Prologue: the two purses ..... 8
A note on the story ..... 9
What a family office can look like ..... 9
The chosen one: the person who starts the family office ..... 10
The prior decision ..... 10
Two starting points ..... 11
The proposed answer: what the barbell is and is not ..... 11
Local knowledge and the limits of foresight ..... 11
Part 1: complexity and family offices ..... 13
Introduction ..... 13
1.1 Complexity made simple: the old measuring problem ..... 13
1.2 What do we mean by complexity? ..... 15
1.3 Bounded vs unbounded systems ..... 17
1.4 Risk vs uncertainty: Knight's distinction ..... 18
1.5 The one-path problem and ergodicity ..... 19
1.6 The progression, with examples ..... 22
1.7 The bridge to heuristics ..... 23
Part 2: heuristics — simple rules for uncertain worlds ..... 24
2.1 What is a heuristic? ..... 24
2.2 What heuristics are for ..... 26
2.3 Why heuristics are necessary for family offices ..... 26
2.5 Nassim Taleb on heuristics ..... 27
2.6 The sizing problem: Kelly and half-Kelly as heuristics ..... 28
2.7 Summary: what heuristics achieve ..... 29
2.8 The bridge to the barbell ..... 29
Part 3: the barbell and supporting heuristics ..... 30
Level 1: the barbell — the shape, the overarching heuristic ..... 30

Level 2: supporting heuristics — the practical habits that keep the barbell honest ..... 32
The argument for seeing them together ..... 34
What this paper does not claim ..... 35
Conclusion: a framework, not a prescription ..... 35
Appendix: the status of the heuristics ..... 35
Contact information ..... 37
Footnote references ..... 38

# INTRODUCTION: WHAT THIS GUIDE OFFERS

Family offices, chief investment officers, limited partners and related professionals all face capital-allocation choices.

The list of places to put money looks long: stocks, bonds, private equity, venture funds, permanent capital vehicles, direct companies, real estate and credit. The labels may differ, but the underlying environment does not stand still. Interest rates change. Liquidity changes. Managers change behaviour. Founders change course. Buyers change what they want. Lenders tighten or loosen terms. Competitors react. New information can alter decisions. These changes can create gains or losses, improve or worsen conditions, and produce consequences—anticipated or not.

This is the setting in which capital operates.

One of the significant issues, therefore, is not simply the investment product. It is the environment in which the product has to perform.

A family office can choose from many investments, but it cannot choose its future sequence of events. It has one balance sheet. It will experience one sequence of capital calls, distributions, gains, losses, taxes, liquidity requirements and family decisions. It cannot live through every possible sequence and then receive the average outcome.

This creates a simple distinction.

An average can describe what happened across a group. A family office has to deal with what happens to it.

That distinction matters when the environment contains changing actors, unequal outcomes, feedback and events whose order can affect what happens next. These are features associated with complex systems. They make prediction more difficult and make the consequences of being wrong more important.

The question is therefore not only: "What might this investment return?"

It is also: "What happens to the family if the outcome is different from what was expected?"

This guide takes that question as its starting point.

It presents these complementary lenses:

1. ***The barbell strategy*** — a capital-allocation shape that places greater weight on capital intended to preserve the family's ability to operate, while reserving a smaller portion for opportunities with substantial upside.
2. ***Complexity*** — a way of understanding why changing actors, unequal outcomes, feedback and sequence can make the future difficult to calculate.
3. ***Uncertainty*** — a way of understanding why probabilities cannot always be assigned, and why some risks cannot be priced in advance.
4. 

The barbell is not presented as the "right" strategy. Nor does the guide argue that traditional portfolio construction is wrong. It presents the barbell as one heuristic for thinking about capital when prediction is limited: preserve the ability to act, while retaining exposure to opportunities that may produce disproportionately large outcomes.

The supporting heuristics that follow are derived from that same starting point. They are intended to make the implications easier to see and discuss, not to replace the judgement of the family and its advisers.

The guide uses simple arithmetic, ordinary examples and storytelling alongside ideas from complexity, uncertainty and decision-making. The purpose is not to make investing simplistic. It is to make the underlying problem easier to see.

This is a framework for thinking about capital allocation under complexity and uncertainty. It is not investment advice, and it does not prescribe a particular allocation for any family.

# THE AUTHORS

This paper was authored by Tiisetso Maloma, Startup Picnic Venture Studio, and Startup Picnic Academy. Tiisetso Maloma founded both the Studio And The Academy.

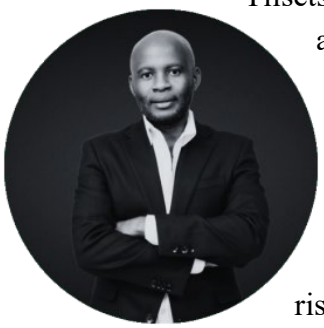

Tiisetso Maloma writes books and tools that help people ideate, innovate and launch. He studies funds, complex systems, innovation, township economies, rhetoric and risk. His books include ***Introducing Ubuntu Stoicism***, ***Quick Innovation MBA***, ***Forget the Business Plan Use This Short Model***, ***The Anxious Entrepreneur***, ***Future of Township Economies***, ***Money is Biological***, and ***90 Days to Create and Launch***.

Startup Picnic Venture Studio builds startups at the intersection of intellectual property, finance, technology and AI. The Studio looks for firms that can scale and last.

Startup Picnic Academy builds short courses and practical tools for founders and trainers. It works with corporates, non-profit organisations and colleges. Delivery is on-site, mobile and online, through classes, workshops, bootcamps, masterclasses and physical games. Courses cover product creation, services, entrepreneurship, innovation and related skills. Some programmes sit against accredited standards; others are unaccredited and built for speed.

The tools used in the Academy's work include the EBC Business Model, the Innovation Design Model and an AI Coach. EBC treats a business as six parts that can be checked: product, marketing, distribution, cash flow, relentlessness and decision-making. The Innovation Design Model is a short sequence: imagine, ask customers, research. The AI Coach is a digital aid that uses those models and examples.

One academy format is the Innovation Teacher Factory: a short fellowship that trains trainers to use those tools with families and youth, including in township settings. That work sits behind the authors' interest in cheap, repeatable training. It is not a finding of this white paper.

# PROLOGUE: THE TWO PURSES

In a village in the Eastern Cape, the King—a traditional leader who had also built a business empire over forty years—called two of his longest-serving helpers. One had cleaned his offices. The other had driven him and, in quiet moments on long roads, told him the truth no one else would. He trusted them more than his board.

He rewarded them the same. Not with a once-off gift that would be spent, but with a trust held for their children and their children's children, so their families would live improved lives forever.

He gave each man R50 million.

The first man was Steve.

Steve took his R50 million and spread it across many careful-looking and sophisticated places. He put R10 million into a balanced unit trust. He put R10 million into a life annuity. He put R10 million into a property fund. He put R10 million into a fund of hedge funds. He put R10 million into a "moderate growth" portfolio at a bank. Each place felt reasonable. Each sent him letters and reports. Each took a fee for the work. On paper, Steve's plan looked balanced. It looked like what a serious man would do. He felt safe.

When the hard times came in—when interest rates moved, when the property market stalled, when the hedge funds locked their doors, when the unit trusts fell with the market—everything Steve had chosen fell at the same time. The letters stopped. The fees had already been paid. Little was left that Steve's children could still use.

The second man was Thabo.

Thabo did something different with his R50 million. He took R35 million and put it in a money market account and short-dated government bonds—places he could reach in days, not months. He took R10 million and bought a small bakery and a fleet of three taxis—businesses he already understood, that already had customers, that already brought in cash every week. He took the remaining R5 million and planted it, year by year, into risky ventures: a small stake in a friend's construction company, a share in a spaza shop chain, a piece of a tech startup in Johannesburg, a loan to a farmer he knew, a minimum amount in a venture capital firm. He kept each of those amounts small enough that even if every one of them failed completely, his family's home would still stand and they would still eat. Between plantings, Thabo kept some of the R35 million aside, untouched, for the next opportunity and the next season.

A friend of both men watched. People called him the Gambler, because he had spent his life around risk—at horse races in Durban, at casinos in Gauteng, in business deals that others thought were mad—and had survived it.

"Why risk any of it at all?" Steve asked him, years later, when it had already gone wrong. "Why not keep everything safe?"

"Because the world does not stand still," the Gambler said. "But listen closely. What Thabo did with his money is what I do with my bets. I never put everything on one horse. I never put everything on one race. I keep most of my money in my bank and bet small on many races across many years. Most of my bets lose. A few win big—covering more than I lost. But because I never bet the whole pocket, I am still here to bet again. Thabo understood the same thing. He did not know which of his small ventures would win. He only knew his family must still be standing to plant again. And he never confused the luck of many gamblers at many tables for the luck he himself would have at his own table."

Years later, the King asked which of the two families could still act—could still help themselves,

could still seize an opportunity, could still recover from a blow.

The Gambler pointed to Thabo. “That one understood the game.”

## A NOTE ON THE STORY

This paper does not claim Thabo was right. It does not claim Steve was wrong. It does not claim the King’s question has one answer. It claims only that the question—which family can still act?—is the question a family office faces, and that the shape of Thabo’s plan is one coherent answer.

What follows is not a derivation from the story. The story does not prove anything. The story shows a way of thinking. The rest of the paper will say that way of thinking plainly, as a set of starting points and simple rules.

# WHAT A FAMILY OFFICE CAN LOOK LIKE

A family office is simply an organised way for a family to manage its wealth and related affairs. It may cover investments, tax, legal work, philanthropy, succession, administration and, in some cases, the operating businesses that created the wealth in the first place. The work can be performed by people employed directly by the family, by external specialists, or through a combination of both. [Deloitte, 2023*][1] [PwC, 2021*][2]

There is no single structure that defines a family office. Some are substantial organisations with investment professionals, accountants, lawyers, administrators and formal governance or investment committees. Others are much smaller: a principal or family member working with one or two trusted advisers. The structure can change as the family’s wealth, needs and complexity change. [Deloitte, 2026*][3] [Family Office Exchange, 2026*][4]

Common arrangements include:

- ***Single-family office***: An organisation established to serve one family. It may employ its own investment, accounting, legal, administrative and other specialist staff. The family generally retains direct control over the structure and the decisions it supports. [Family Office Exchange, 2026*] [J.P. Morgan, 2026*][5]
- ***Multi-family office***: A professional organisation serving several families. Infrastructure, specialist staff and operating costs can be shared across those families, while each family maintains its own arrangements and decision-making. A multi-family office can therefore provide some of the capabilities of a dedicated office without requiring one family to carry the entire cost of building and maintaining the infrastructure. [Family Office Exchange, 2026*] [PwC, 2021*]
- ***Outsourced family office***: The family retains responsibility for major decisions but uses external providers for some or many functions. Tax, legal work, accounting, investment management, reporting, administration and other specialist functions can be performed outside the family office. In practice, many family offices use a mixture of internal and external resources rather than keeping every function in-house. [Deloitte, 2023*] [Deloitte, 2026*]
- ***Virtual family office***: A flexible, technology-enabled version of an outsourced or hybrid model. Rather than maintaining a large permanent office and staff, the family can coordinate a network of external specialists while using technology for reporting, communication, document management and administration. The family can retain control of decisions while much

of the work is performed remotely. [PwC, 2021*] [U.S. Bank, 2026*][6]

- ***Operating-company family office***: The wealth-management function remains closely connected to the business or businesses that created the family's wealth. Some family-office activities may therefore sit within, or alongside, the operating business rather than being separated into a completely independent structure. [PwC, 2021*]

These arrangements are not mutually exclusive. A family might have a small internal team, outsource tax and legal work, use external investment managers and coordinate the different providers through a technology-enabled system. Another family might employ a substantial internal team and outsource only highly specialised functions. Deloitte similarly describes family-office services as a mixture of activities that may be performed in-house, outsourced or delivered through a combination of both. [Deloitte, 2023*]

The important point for this guide is that a family office does not have to be a large building filled with employees. It can be a formal organisation, a small internal team, a multi-family platform, or a coordinated network of people and firms. What matters is the organised way in which the family manages its capital, affairs and long-term interests.

# THE CHOSEN ONE: THE PERSON WHO STARTS THE FAMILY OFFICE

"The Chosen One" is this paper's name for the person who establishes a family office. It is a reference label, not a claim. It gives the discussion a person through whom the decisions involved can be examined.

The Chosen One does not have to begin with substantial wealth. Some establish a family-oriented structure early, while still building wealth, wanting the businesses, investments and assets they accumulate to serve a longer-term family purpose from the outset. Others establish a family office much later, after substantial wealth already exists—following a business sale, an inheritance, investment success, a succession concern, a change in family circumstances, or simply the recognition that the wealth has become large or complex enough to require a different structure. Some understand the idea early. Others arrive at it only after creating significant wealth. In every case, the Chosen One is the person who decides that wealth should be organised with the family and its longer-term interests in mind.

The wealth may have come from entrepreneurship, employment, investments, an inheritance, a business sale, other means, or even luck—the lottery, for example.

The Chosen One may also still be creating further wealth. They may still be running the business that created their wealth. They may be pursuing another venture, backing a venture fund, establishing a permanent-capital vehicle, investing directly in businesses, or pursuing another opportunity. Wealth creation may therefore continue alongside wealth preservation.

## THE PRIOR DECISION

This creates a decision that comes before everything else in this paper:

How much of the wealth should become family-office capital?

The answer belongs to the Chosen One. The whole fortune does not necessarily have to enter the family office.

Consider a Chosen One with R100 million in wealth. One possibility is that the full R100 million

becomes family-office capital. Another is that R40 million is judged sufficient for the family office, with the remaining R60 million kept outside it. That R60 million might remain in the business that created the wealth. It might fund another venture, back a venture fund, sit in a permanent-capital vehicle, or be deployed into another opportunity. The Chosen One may be willing to lose it in the pursuit of creating more wealth. If it succeeds, some of the resulting wealth may eventually enter the family office—or it may fund another venture, support philanthropy, or serve another purpose.

This guide does not attempt to determine that amount. It begins after that decision.

# TWO STARTING POINTS

## Scenario 1: R100 million enters the family office.

The Chosen One decides that the entire R100 million is now family-office capital. That capital has a preservation function, though preservation does not mean every rand must sit in cash or low-risk assets. The barbell proposed in this guide is then applied to the R100 million.

## Scenario 2: R40 million enters the family office.

The Chosen One decides that R40 million is sufficient to form the family's long-term capital base. The remaining R60 million stays outside the family office and remains available for continued wealth creation. The barbell is then applied to the R40 million.

The R60 million is not treated as family-office capital for the purposes of this guide. It may succeed, fail, or produce an outcome somewhere between the two. If it succeeds, some of the resulting wealth may later be transferred into the family office, increasing the capital available for the family's long-term purposes.

The two scenarios begin differently, but the question this paper addresses is the same: once capital has been designated as family-office capital, how can it be thought about under complexity and uncertainty?

# THE PROPOSED ANSWER: WHAT THE BARBELL IS AND IS NOT

The answer proposed here is the barbell. Within the family office, the preservation side remains important. A smaller portion can be exposed to higher-risk opportunities, but the purpose of that exposure is not to endanger the capital base. It is to retain participation in upside while keeping the larger base able to stand.

The barbell is therefore not a prescription for how much wealth the Chosen One should place in a family office. That decision comes first. The barbell is a proposed way of thinking about the capital once it is there. The same logic applies whether the family office contains R100 million, R40 million, or another amount entirely. The Chosen One remains the decision-maker.

# LOCAL KNOWLEDGE AND THE LIMITS OF FORESIGHT

If the Chosen One is still running the venture that created the wealth, they may possess valuable local knowledge of its customers, suppliers, operations, margins and cash cycle. That knowledge can improve understanding of the present situation. It does not make the future predictable. Customers can change. Competitors can respond. Costs can rise. Technology can shift. New opportunities can appear. The venture remains exposed to an environment that changes.

Later generations may inherit the family-office capital, the structures around it, and—where wealth creation continues—businesses and investment vehicles outside the family office. The underlying caution remains the same: the house must still stand after a bad sequence.

This paper therefore starts after the Chosen One has made the first decision. It does not prescribe how much wealth should enter the family office. It asks how the capital designated for that purpose can be thought about when the future cannot be calculated with confidence.

# PART 1: COMPLEXITY AND FAMILY OFFICES

## INTRODUCTION

This part of the paper does one thing. It shows why the future is hard to calculate.

It does not begin with a definition. It begins with a problem that anyone who has ever counted anything already understands. Counting. Averaging. A doorway that is tall enough. A herd of cattle divided between households. These are ordinary acts. They are where the trouble starts.

That trouble has a name. Complexity. But the name is not the point. The point is what complexity does to a family office trying to decide where its capital should sit.

### The word and the thing

In everyday use, "complex" means "difficult to understand." A watch is complex. An aeroplane is complex. A family business is complex. A derivatives portfolio is complex. The word gets applied to anything that feels hard.

That usage is not wrong, but it points at the wrong thing. It describes how an object feels to us, not how it behaves. A watch has thousands of parts and is still fully predictable. Difficulty and complexity are not the same problem.

Complexity science is itself a developing field. There is still no single definition that all its researchers accept. Some have tried. Herbert Simon offered one in 1962 that is still used. Others have argued that the features commonly attached to complexity—non-linearity, feedback, emergence—are neither necessary nor sufficient. Even the boundary between "complicated" and "complex" has never been sharp. A watch has thousands of parts, but its behaviour is fully determined. A worm's brain has a few hundred neurons, and it does things no one can predict from the neurons alone.

It has been said that complexity is the tombstone of definitions.

This chapter does not try to solve that problem. It does not offer a formal definition that settles the debate. It does something narrower and, for a family office, more useful. It shows how the ordinary act of measuring breaks down when the thing being measured becomes variable, scalable and dependent on other things. It then shows what the formal literature says about systems with those properties. And it lands on the one consequence that matters most for a family office: it lives one path, not the average of every path.

## 1.1 COMPLEXITY MADE SIMPLE: THE OLD MEASURING PROBLEM

This segment is written in deliberately simple and, in places, illustrative terms. It is not intended as a history of mathematics or a formal definition of complexity. It is a way of showing how a very ordinary problem begins: we can count things quite easily, but the meaning of the number becomes less certain as the thing being counted becomes more variable, scalable and dependent on other things, i.e. it is complex.

The trouble with complexity often starts with the simplest act: counting.

Long before mathematics became a discipline, people could compare quantities. Which household had more cattle. Which had less grain. Perhaps a herder cut notches into wood, or set pebbles aside: when a cow left the kraal, a pebble moved to one side; when the cow returned, it moved back. If a pebble remained on the wrong side, something was missing. Counting in this case, as in many cases, was not just an abstract exercise. It was a practical

way of keeping track of things that could be seen, touched and had a cost.

The invention of numerals made this faster. Writing "5" instead of five marks. Writing "1,000" instead of a thousand strokes. Symbols allowed people to handle larger numbers and more complex arithmetic. But the symbols themselves did not change the underlying problem: the meaning of a number depends on what is being measured, and on how the things being measured relate to one another.

Builders understood this. A door must be tall enough for a person to pass through. Human height varies, but it is biologically bounded. Almost nobody is three metres tall. So a carpenter does not need the exact height of every future user. A doorway around two metres high, with appropriate margin, will accommodate most people. A table around eighty centimetres high—roughly half of average human height—will be usable by many adults. The average height does not correspond to any particular person, but for the purpose of building a doorway, it is good enough. And the half of average human height for a table is inconsequential.

This is an important property: when things are bounded and the cost of being approximately wrong is manageable, the average can be used.

The problem begins when the boundaries start to disappear.

Three households own 2 cattle, 3 cattle and 1 cow. Average: 2. Another village: one household owns 6, the other two own none. Average: still 2. The arithmetic is identical. The reality is not. The average has not become incorrect. It has simply stopped telling the whole story.

When numbers lose their natural upper limit, the problem gets worse. Three people earn R10,000, R20,000 and R30,000 a month. Average: R20,000. That number roughly describes the group. Replace the third income with R1,000,000. Average: R343,333. The arithmetic is still correct. But R343,333 no longer describes what any of the three people actually earns. One observation has changed everything.

A person's height cannot grow from two metres to two hundred metres. A company can. An investment can. One founder can create an outcome larger than the combined outcomes of thousands of ordinary businesses. Economic quantities can be highly unequal and, in practical terms, lack the kind of natural ceiling that makes physical averages intuitive [Taleb, 2007][7].

This is why the language of "typical" outcomes requires care in finance. Consider ten investments, each of the same size—R1 million allocated to each. Nine return +10% (each turning R1 million into R1.1 million). One returns +1,000% (turning R1 million into R11 million). The arithmetic average of the ten percentage returns is 109%. That figure is a correct description of the ten observed returns. Yet nine of the ten investments produced nothing close to 109%.

If the question is "What is the average of these outcomes?", the answer is 109%.

If the question is "What should a randomly selected investment return?", the answer is far less obvious—most of the time one would simply receive the ordinary +10%.

If the question is "What will happen to this family office?", the average becomes still less useful. The office does not own the average. It owns particular investments and experiences their particular sequence of outcomes.

Now contrast this with a situation in which the investments are of unequal size. Suppose the family office has allocated capital as follows:

- R10 million into Investment A, returning +40% (a profit of R4 million)

- R5 million into Investment B, returning –60% (a loss of R3 million)
- R3 million into Investment C, returning –50% (a loss of R1.5 million)
- R2 million into Investment D, returning –25% (a loss of R0.5 million)

The simple arithmetic average of the four percentage returns is –23.75%. Even though the single highest return is a solid +40%, the larger losses on the other positions, weighted by the capital actually invested, produce an overall portfolio loss of R1 million on the R20 million deployed.

Once again, the office does not own any average. It owns the specific investments in the specific amounts it chose, and it lives with the concrete sequence of gains and losses those particular positions deliver. Looking at the numbers available today, it is not clear what the future will be.

This is the beginning of complexity.

## 1.2 WHAT DO WE MEAN BY COMPLEXITY?

If the average cannot carry the decision, then something else must. Yet it is not immediately obvious what that something else should be, or how it should carry the weight of choice.

Some variables are known and estimatable. These can be acted upon, within limits. The inflation rate over the next quarter is known to exist. Central banks publish targets. Economists produce forecasts. A family office can examine the range and decide, with an appropriate margin, how much of its safer capital should sit in short-dated instruments. It cannot know the exact figure, but it can still act.

Some variables are known but not estimatable. These can be named, watched and prepared for, yet they cannot be reduced to a reliable number. The behaviour of a competitor is known to matter. The competitor exists; its next move will affect the business. No formula, however, produces the competitor's next price change, its next hiring decision, or whether it will cut margins to gain share. The variable is visible; it is not calculable.

And some variables are not yet known at all. These cannot be named, watched or prepared for, because no one yet knows they exist. A technology that has not been invented, a regulation that has not been drafted, a buyer that has not entered the market, a founder's health event that has not occurred—these are unknown not merely in value but in existence. No amount of historical data can reveal a variable that has not yet appeared.

Even the known and estimatable variables fail to disclose how they will interact with everything else. Two brief examples illustrate the point.

***First,*** interest rates and a property portfolio. The family office holds a portfolio of rental properties. Interest rates are a known and estimatable variable; a reasonable range can be forecast. On its own the variable appears manageable. Yet rates interact with the cost of existing debt, the affordability of tenants, the willingness of buyers to pay current valuations, and the pipeline of new supply. A rise that looks modest in isolation can, when combined with an already stretched tenant base and additional supply, produce a vacancy problem no single forecast anticipated. The variable was measured accurately; its interaction was not.

***Second***, a key customer and a founder-led business. The family office invests in a company whose revenue depends on one large customer. The customer's existence, the contract terms and the renewal date are all known and, within limits, estimatable. But the customer is simultaneously reacting to its own market, its own competitors and its own internal politics—none of which the family office controls or can fully observe. If the

customer's business softens, it may delay renewal, demand a price cut or switch suppliers. None of these outcomes is written into the contract. The variable was known; its interaction with the wider system was not.

Both examples lead to the same conclusion: knowing a variable, and even estimating it accurately, does not reveal how it will behave once it meets the other variables in the system. A model can assign a number to each known factor; the number does not make the interaction predictable.

The moment capital is actually deployed, two further problems appear at once: how much capital goes into each position, and how those positions sit together.

Sizing and allocation. Neither problem can be solved by a single average. Neither can be solved by the known variables alone. The first asks how much of the safer capital should fund a position of a given size. The second asks what happens when several positions, each reasonable in isolation, respond to the same shock at the same time.

This is the point at which complexity becomes a practical problem rather than an abstract one.

Begin with scale. A person's height cannot grow from two metres to two hundred. A company can grow from a small enterprise into a billion-rand business. An investment can lose almost everything while another returns many times the original capital. One founder can produce an outcome larger than the combined results of thousands of ordinary businesses. Economic quantities are highly unequal and, in practical terms, lack the natural ceiling that makes physical averages intuitive [Taleb, 2007]. In Taleb's terms, economic life belongs to Extremistan, where "exceptions can be everything" and a single observation can dominate the rest. The first property of the systems family offices confront is therefore this: things scale without a hard upper limit.

Next consider the actors. Even the variables that can be observed belong to agents who respond strategically. W. Brian Arthur put the matter directly: "economic elements (human agents) react with strategy and foresight by considering outcomes that might result as a consequence of behavior they might undertake." Founders react to competitors; competitors react to founders; investors react to returns; banks alter lending conditions. The thing being measured participates in the process that generates the next observation. Arthur's formulation captures the loop: "As the elements react, the aggregate changes; as the aggregate changes, elements react anew." Uncertainty is not confined to the variables themselves; the actors behind them continually revise their behaviour—including in response to the family office's own decisions [Arthur, 1999][8].

This produces the harder problem. The behaviour of a single actor can sometimes be predicted, or at least bounded. What cannot be predicted is how the actors will interact. Melanie Mitchell's work on complex systems emphasises precisely this character: "a system with large numbers of interacting components, in which the components are relatively simple compared with the system as a whole, in which there is no central control or global communication among the components, and in which the interactions among the components give rise to complex behavior." The outcome emerges from the interaction, not from any one participant's plan [Mitchell, 2009][9].

Complexity is pervasive in investing for exactly these reasons [Hsu & West, 2025*][10]. New products, new processes, financial engineering and innovation continually add layers to an already intricate system [Holland, 1992][11]. Over time the infrastructure designed to manage that complexity often expands—more specialists, more processes, more coordination [Mitchell, 2009]. For family offices the difficulty is not confined to investments.

Families evolve organically; with each generation new personalities and opinions appear, complicating matters through conflicting agendas and differing priorities [Coyle, 2025*][12]. Complexity consumes both money and attention [Cambridge Associates, 2023*][13]. It also sustains the belief that complex problems require complex solutions [Hsu & West, 2025*]. As Edsger W. Dijkstra observed, "to make matters worse: complexity sells better" [Dijkstra, 1989][14].

## So what, precisely, is complexity?

In everyday language the word often means little more than "I do not quite understand it." A mechanical clock is structurally elaborate, yet its behaviour remains predictable. Complexity concerns a different class of systems: those whose behaviour cannot be derived simply by analysing their separate components, but emerges from the interactions among them. Melanie Mitchell said "a formal definition of 'complex system' is impossible at this time"—because the key terms (adaptive, emergent, intelligent) themselves lack precise meanings [Mitchell, 2009][9]. Stephen Wolfram put the difficulty more bluntly: "Formal definitions can get all tied up in knots—just like formal definitions of almost anything fundamental: life, energy, mathematics, etc." [Wolfram, 2008*][15].

Herbert Simon, in *The Architecture of Complexity* (1962), offered a definition that remains useful: a complex system is "one made up of a large number of parts that interact in a non-simple way." In such systems "the whole is more than the sum of the parts, not in an ultimate, metaphysical sense, but in the important pragmatic sense that, given the properties of the parts and the laws of their interaction, it is not a trivial matter to infer the properties of the whole" [Simon, 1962][16].

Complexity theory draws attention to certain recurring features of these systems [Holland, 1992] [Mitchell, 2009]:

- Non-linearity — small inputs can produce effects that are not proportional [Arthur, 1999]
- Adaptation — agents change their behaviour in response to events [Holland, 1992]
- Emergence — outcomes arise from interaction rather than central planning [Mitchell, 2009]
- Path dependence — where a system is now depends on the path it has already taken [Arthur, 1989][17] [David, 1985][18]

These features imply that prediction is inherently limited in complex environments [Knight, 1921][19]. The question complexity theory therefore raises is not "How do we predict better?" but "How do we make decisions when prediction is limited?" [Gigerenzer & Goldstein, 1996][20] [Gigerenzer, 2008][21].

# 1.3 BOUNDED VS UNBOUNDED SYSTEMS

This section gives the problem its formal vocabulary.

There is a useful way to think about when averages work and when they do not. It comes from the distinction between bounded and unbounded systems.

***Bounded systems (Mediocristan).*** Bounded systems are those where outcomes have a natural limit—a ceiling and a floor. Human height works for building doorways because it is biologically constrained. Aeroplanes are bounded by the laws of physics and the limits of materials. The calendar is bounded—24/7, 365 or 366 days. These systems are predictable because the rules are stable and the outcomes are bounded [Knight, 1921].

Nassim Nicholas Taleb calls this world Mediocristan [Taleb, 2007]. In Mediocristan:

- Exceptions occur but do not carry large consequences
- Add the heaviest person on the planet to a sample of 1,000 people. The total weight barely changes
- Individual observations do not dominate the average
- The average works. Probability works. You can build models with confidence.

***Unbounded systems (Extremistan).*** Economic systems are different. Unbounded systems are those where outcomes can scale without a hard upper limit: wealth is unbounded, investment returns are unbounded, company valuations are unbounded. In unbounded systems:

- One observation can dominate the average
- One founder can create an outcome larger than thousands of others combined
- One investment can wipe out a portfolio or multiply it many times over

Taleb calls this world Extremistan [Taleb, 2007][7]. In Extremistan:

- Large deviations occur and they are consequential
- Add Bill Gates—or any dollar billionaire—to a sample of 1,000 people. The wealth jumps by a factor of hundreds of thousands
- The average becomes meaningless. One observation changes everything
- The median might be zero. The system does not behave like a stable physical system

***Adaptive systems.*** Economic systems are also adaptive. A table does not change because someone measures it. People do. Founders respond to competitors. Investors respond to returns. Banks change lending conditions. Consumers change preferences. The thing being measured therefore participates in the process that produces the next observation [Arthur, 1999].

Take two entrepreneurs who start with similar businesses. One obtains a major customer early. The additional cash allows that entrepreneur to hire staff, improve the product, acquire more customers and obtain further financing. The initial difference creates later differences. The same input—two similar businesses—produces different outputs depending on the order in which events arrive [Arthur, 1989] [Arthur, 1999].

That is the shape of the problem. Outcomes scale without a bound. The actors adapt. The sequence matters. And the system keeps changing while the family office is still inside it [Holland, 1992] [Mitchell, 2009].

# 1.4 RISK VS UNCERTAINTY: KNIGHT'S DISTINCTION

Some situations let you assign useful odds. Many of the decisions a family office makes do not.

In 1921, Frank Knight drew a clear line:

- Risk can be measured with probabilities.
- Uncertainty cannot.

As Knight put it: there is a basic difference between the reward for taking a known risk and the reward for taking a risk whose size is itself unknown. A known risk can be turned into something close to certainty. True uncertainty cannot be measured at all [Knight, 1921].

## Risk: the measurable kind

Risk exists when you know the possible outcomes and can estimate their probabilities—either from the situation itself or from past data.

A clear example is funeral insurance. The insurer looks at historical death rates and the typical death rates in a country. From these past patterns it infers that, in a large group of policyholders, only a certain number are likely to die within a given period, and it prices the premiums accordingly. It

can also estimate the point at which an unusually high number of claims would threaten the company, and it can buy reinsurance to protect against that extreme outcome. The premium therefore reflects these measurable, statistically estimated risks. The outcome for any single person remains uncertain, but the overall distribution for the group can be estimated and managed with reasonable confidence.

Knight's point was that this kind of measurable uncertainty is not really uncertainty in the economic sense. It can be turned into a cost (an insurance premium) and absorbed into the price.

## Uncertainty: the kind that cannot be measured

Uncertainty is different. It arises when you do not fully know the range of possible outcomes and cannot assign probabilities with any confidence.

A clear example is a startup selling something entirely novel—such as an AI product that has never been tested in the real world. There is no reliable distribution of outcomes for "this specific team, this specific product, this specific market, at this specific moment." You cannot look at enough similar past cases, run the experiment a thousand times, or average the results. You cannot insure against it or price it with confidence. The situation is unique.

***Other examples:***

- The risk that AI leads to human extinction (we have never built such a system before).
- The commercial success of electric vehicles in 2010 (you cannot repeat the invention enough times to get statistical odds).
- The economic effects of a pandemic.
- The invention of the internet, the rise of Bitcoin, or the early days of Covid.

These are risks that cannot be priced in advance. They were uncertainties that had to be lived through.

## Where private capital sits

A lot of private capital sits closer to uncertainty than to risk.

In risk, the range of outcomes and their probabilities are known or can be estimated. In uncertainty, neither is fully known.

A venture capital fund cannot tell you the exact probability that a particular company will succeed. It can cite historical base rates for startups in a sector, but a base rate describes a group, not this unique company. There is no large set of identical cases to average over.

Knight's insight was that this is not a flaw in the analysis—it is a feature of the real economy. Profit (the extra return entrepreneurs earn) exists precisely because uncertainty cannot be insured away. If it could be priced, competition would eliminate the extra return. The entrepreneur is paid for bearing what cannot be measured.

## Why this matters for a family office

This distinction is not minor. It changes how decisions should be made.

When you face risk, you can calculate. You can size positions with formulas, diversify across many independent risks, and rely on the law of large numbers. You can build and stress-test models.

When you face uncertainty, you cannot calculate. You do not know the distribution. You cannot treat outcomes as independent draws—they interact, adapt and depend on the path taken. This paper infers you need heuristics that protect you from what you cannot measure.

# 1.5 THE ONE-PATH PROBLEM AND ERGODICITY

This is complexity's most important consequence for a family office.

***The one-path problem.*** A statistical study can look at thousands of investments in a certain class or category at once and work out the average return across all of them. But as we have established, an average can hide losses, and no one knows in advance which investments will be the right ones. A family office cannot hold all of those thousands, and it cannot know in advance which ones will succeed. It has one balance sheet and one sequence of capital calls, distributions, exits, losses, taxes and family requirements. It cannot experience every possible investment path and then receive the average. It experiences its own.

***Ergodicity and non-ergodicity.*** This problem has a formal name: ergodicity.

An ergodic system is one where the time average (e.g. what one person experiences over time) equals the ensemble average (the average across many parallel cases). If you track one system over time and average the results, and you average across many parallel systems at the same moment, you get the same number. When this holds, a probability description can substitute for a time description [Peters, 2019][22].

## Ergodic examples: when the two averages agree

- ***Repeated fair coin tosses or die rolls (additive outcomes):*** the long-run frequency of heads in one sequence converges to ½, matching the average across many independent coins at any fixed moment.
- ***An ideal gas in thermal equilibrium:*** the time average of a molecule's kinetic energy equals the ensemble average across many molecules at one instant. In these cases a single long path is representative of the whole, so expected-value language is safe.
- ***A clock.*** Watch one clock for a thousand hours. Count how often the minute hand passes the twelve. Roughly once every sixty minutes. Now look at a thousand clocks, all at the same moment, and count how many have their minute hand near the twelve. Roughly one in sixty. The time average matches the ensemble average. The clock's behaviour over time is indistinguishable from the behaviour of many clocks at one instant. This is what ergodicity looks like.
- ***A simple interest account.*** Suppose a family office places R1 million in a fixed deposit paying 10% simple interest per year. After 10 years, it has R2 million. If the office had instead placed R1 million each into 10 separate identical accounts, and averaged the outcomes at year 10, the average would also be R2 million. The time average equals the ensemble average. The system is ergodic. The average is a valid guide to what the office will actually experience.

## Non-ergodic examples: when the two averages diverge

- ***The coin toss.*** Start with $100. Flip a fair coin. Heads: wealth increases by 50%. Tails: wealth decreases by 40%. Now imagine 10 people, each starting with $100 and playing once. Five get Heads and end with $150 each. Five get Tails and end with $60 each. Together, they have $1,050, so the average is $105 per person—a 5% gain (ensemble average). But follow one person over two rounds. If they get Heads and then Tails, their money goes from $100 to $150 and then to $90. If they get Tails and then Heads, it goes from $100 to $60 and then to $90. Either way, one Head and one Tail leaves the person with $90. They started with $100 and now have $90. The group average went up, while the individual path

went down. This is the basic idea of non-ergodicity: the average across many people can differ from what one person experiences over time [Peters & Gell-Mann, 2016][23] [Peters, 2019].

- The concentrated position. Suppose a family office allocates R1 million to a single early-stage company. Statistical studies of this class of investments might show that, across many such companies, a small number return 1,000%—turning R1 million into R11 million—while a larger number return nothing. The average across that whole group might therefore be positive, even though most of the companies in it returned nothing. But the family office does not hold the whole group. It holds one. If the company fails, the R1 million is gone. The office cannot average its way out because it cannot experience the group. It experiences one outcome. The group average says the investment is attractive. The time average for this particular office may say it lost everything. The system is non-ergodic.

## Why this matters for a family office

The coin-toss example shows the basic problem. The average result across many people is not necessarily the result experienced by one person over time.

Peters and Gell-Mann call this distinction the difference between an ensemble average and a time average [Peters & Gell-Mann, 2016]. The ensemble asks: What is the average result across many possible cases? The time average asks: What happens to one person as the process unfolds over time?

For a family office, the second question matters. A family office does not own hundreds of parallel versions of its balance sheet. It has one balance sheet. It experiences one sequence of investments, gains, losses, capital calls, distributions, taxes and family needs. Taken together, these have to leave enough capital for the family office to remain financially viable over time. The family office needs to remain able to operate, meet obligations and preserve its capital base through time.

This means that a positive expected return does not guarantee a positive outcome for a particular family office. The office may hold many investments, but it lives through only one sequence of events—one path through time. It may experience a loss followed by a gain, but the timing, size and order of those outcomes are not known in advance. It cannot choose a different sequence simply because that sequence would have produced a better result. If a loss is large enough to impair the office's ability to continue, a later opportunity may arrive too late. This is why the sequence matters. Peters argues that wealth is a particularly important example because wealth changes through time. Each gain or loss changes the amount of capital available for what comes next. R100 million followed by a 50% loss leaves R50 million. A later 50% gain takes R50 million to R75 million, not back to R100 million. The family office therefore has to remain able to operate through the sequence it actually experiences [Peters, 2019].

The practical lesson is simple: An average can describe many possible outcomes. A family office has to survive its own outcome.

This does not mean that probabilities, expected returns or investment models are entirely useless. It means they should not be treated as though they describe with certainty what one balance sheet will experience.

That is the one-path problem.

The family office does not hold the average. It holds the capital that actually survives its own sequence.

# 1.6 THE PROGRESSION, WITH EXAMPLES

Each stage asks a harder question than the one before. The examples below show what each stage looks like in practice—and where the answers start to break down.

| ***Stage*** | ***Question*** | ***Example*** | ***Does the answer hold?*** |
|---|---|---|---|
| ***Counting*** | How many? | A herder counts cattle leaving the kraal and returning. Each cow is one pebble. If a pebble is left over, a cow is missing. | Yes. The count is exact. |
| ***Averaging*** | What is typical across a group? | A carpenter builds a table. Human height varies, but it is bounded. A table at roughly half of average human height—about eighty centimetres—suits most adults. | Yes—because height is bounded and the cost of being slightly wrong is small. |
| ***Probability*** | What outcomes and likelihoods can be specified? | A funeral insurer prices premiums from historical death rates. It cannot say who will die, but across a large group it can estimate how many will. | Yes—because the group is large and the distribution is stable. |
| ***Complexity*** | What happens when many changing actors interact, when outcomes become highly unequal, and when the sequence matters? | A family office backs a founder-led business whose revenue depends on three large clients. One client is acquired and pauses its orders. A second demands a price cut. The third delays renewal while it renegotiates with its own suppliers. Each client was known and each contract was signed; how the three would move together was not. | No. The clients were known; their interaction was not. |
| ***Non-ergodicity*** | What does the group average mean for one decision-maker living one path? | A study of 100 early-stage companies shows a positive average return, because 5 winners returned 1,000% each. A family office invests in one company. It is one of the 60 that returned nothing. The average was positive; its outcome was zero. | No. The group average was never available to the one office. |

### *How to read the table:*

The first two rows can be answered with confidence. The third can be answered if the group is large and stable. The last two cannot be answered reliably—not because the arithmetic is wrong, but because the thing being measured does not behave like the things in the rows above.

At that point, arithmetic has not failed. Probability has not failed. The problem is that neither necessarily describes the path that one family office will live.

## 1.7 THE BRIDGE TO HEURISTICS

Count what can be counted. Estimate what can reasonably be estimated. Where the future cannot be reliably calculated, use simple rules that preserve the ability to act.

That is the bridge to the barbell.

# PART 2: HEURISTICS — SIMPLE RULES FOR UNCERTAIN WORLDS

## 2.1 WHAT IS A HEURISTIC?

The word heuristic comes from the Greek ***heuriskein***, meaning "to find" or "to discover." It is the same root as Archimedes' famous exclamation "Eureka!"—"I have found it!" A heuristic is a finding tool. It is a rule of thumb, a mental shortcut, or a practical method that helps a person find a solution when a complete answer is not available.

Heuristics are emergent in the sense that they are not designed in advance. They develop from repeated interaction with an environment: the decision-maker recognises what works, discards what does not, and stabilises the rule through repetition. This process is described in the fast-and-frugal literature, where cue orderings and rule use can arise from individual learning rather than innate design or formal instruction [Gigerenzer, Hoffrage & Goldstein, 2008][24].

In academia, heuristics have been studied across disciplines:

| Field | Contribution |
|---|---|
| Psychology | Kahneman and Tversky showed that people use heuristics to make judgments under uncertainty [Tversky & Kahneman, 1974][25]. These mental shortcuts are often useful but can also lead to systematic biases. |
| Computer science | Heuristics are used in algorithms to find good-enough solutions when optimal solutions are too expensive to compute [Simon, 1990][26]. |
| Economics | Simon introduced the concept of bounded rationality—the idea that human decision-making is constrained by limited information, time and cognitive capacity [Simon, 1955][27]. Heuristics are the tools we use to work within these bounds. |
| Decision science | Gigerenzer and colleagues developed the "fast and frugal" approach, arguing that simple heuristics can be more accurate than complex models in many real-world environments [Gigerenzer & Goldstein, 1996] [Gigerenzer, 2008]. |

## 2.2 WHAT HEURISTICS ARE FOR

One view, developed most fully by Gerd Gigerenzer and colleagues in the "fast and frugal" tradition, holds that heuristics are not a lesser form of calculation. On this view, they are not a substitute for calculation when calculation is possible. They are a substitute when calculation is not possible—or when it would take too long, cost too much, or produce a false sense of precision [Gigerenzer, 2008].

Within that tradition, heuristics are understood to serve several purposes:

- Speed — they allow decisions to be made quickly when time is limited.
- Frugality — they use only the information that is necessary, ignoring what is not.
- Robustness — they can work across different environments without needing to be recalibrated.
- Protection — they can reduce exposure to catastrophic outcomes by setting boundaries.

The "fast and frugal" approach does not claim that heuristics are crude or inferior. It argues that they are matched to the environment in which they are used, and that a simple rule can outperform a complex model when the environment is uncertain and the data are limited [Gigerenzer, 2008] [Gigerenzer, Hoffrage & Goldstein, 2008].

- The gaze heuristic is one example from that literature. A cricket fielder running to catch a high ball does not solve the physics of a parabola. The fielder fixes their eyes on the ball and runs so that the angle of their gaze stays constant. The catch happens without the mathematics ever entering the mind. The same rule has been described in fighter-pilot interception, where the underlying equations are too slow to solve in real time [Gigerenzer, 2008].
- Gigerenzer also cites the 2009 landing of US Airways Flight 1549. When both engines failed after a bird strike, the pilots did not calculate whether they could glide back to LaGuardia. They fixed their gaze on the airport tower. If the tower rose in the windshield, they could not reach it. The tower rose. They landed in the Hudson River instead. Everyone survived [Gigerenzer, 2008].

For a family office facing uncertainty, a similar logic may apply: a simple rule that ignores some inferences and keeps the house standing after negative movements.

## 2.3 WHY HEURISTICS ARE NECESSARY FOR FAMILY OFFICES

Heuristics are necessary for family offices for three reasons:

### a) Uncertainty cannot be calculated

As Knight distinguished, risk can be priced; uncertainty cannot [Knight, 1921]. Most private capital decisions sit closer to uncertainty than risk. The future path of a portfolio cannot be reduced to a probability distribution. When probabilities are not available, heuristics provide a way to act without pretending to know what cannot be known.

### b) The one-path problem

A family office cannot rely on the average. It cannot hold every investment in a class. It holds a few. So it never earns the class average—it earns whatever its few holdings do.

It cannot average its experience across every other office. It lives one sequence of outcomes. Even if it held the whole class, it would still live one path through time.
A heuristic that preserves the ability to survive that sequence is more valuable than a model that describes the average experience of many offices.

### c) Complexity grows faster than understanding

As shown in this paper, complex systems adapt, produce non-linear outcomes and exhibit path dependence. New products, new managers, new structures and new risks continue to emerge [Holland, 1992] [Mitchell, 2009]. A model built yesterday may not describe today. A heuristic that works across changing environments may be more robust than a model that must be constantly updated.

Dijkstra once observed that "complexity sells better" [Dijkstra, 1989]. It has also been argued that complexity carries real costs—in money and in attention [Cambridge Associates, 2023*]. One view holds that the financial industry has an incentive to make investing appear more complex than it is, because complexity can justify fees. On this view, more products, more strategies and more layers of analysis can create the appearance of sophistication while adding little value.

The implication is not that sophisticated analysis is bad. It is that more calculation does not automatically produce better decisions. As Gigerenzer argues, a decision rule should fit the environment in which it is being used [Gigerenzer, 2008].

For a family office, this creates a practical distinction. The question is not always, "Can we build a more complicated model?" It may instead be, "What simple rule keeps us from making a decision that could permanently impair the family?" That is not anti-mathematical. It is an attempt to match the method of decision-making to the limits of what can actually be known.

## 2.5 NASSIM TALEB ON HEURISTICS

Nassim Nicholas Taleb has written extensively on the role of heuristics in navigating uncertainty. His work emphasises the distinction between probabilistic knowledge and heuristic knowledge:

- Probabilistic knowledge — what you know about the distribution of outcomes. This works in Mediocristan, where outcomes are bounded and probabilities are stable.
- Heuristic knowledge — rules that protect you when you don't know the distribution. This works in Extremistan, where outcomes are unbounded and probabilities cannot be reliably estimated.

Taleb argues that heuristics are not inferior to probabilistic models. They are the appropriate tool when the environment is complex and uncertain [Taleb, 2007] [Taleb, 2012][28].

Key heuristics Taleb advocates include:

| Heuristic | Meaning |
|---|---|
| Via negativa | Know what to avoid, not just what to pursue. The most important decisions are often about eliminating downside, not maximising upside. |
| Barbell strategy | Place most resources in safe, liquid assets and a small portion in high-risk, high-upside bets. Avoid the middle, which offers neither |

| Heuristic | Meaning |
|---|---|
| | safety nor significant upside. |
| Redundancy | Hold more than you think you need. The cost of redundancy is lower than the cost of being caught without resources in a crisis. |
| Skin in the game | Those who make decisions should bear the consequences. This aligns incentives and prevents reckless risk-taking. |

Taleb's barbell strategy is the central heuristic of this paper. It is explored in detail in Part 3.

# 2.6 THE SIZING PROBLEM: KELLY AND HALF-KELLY AS HEURISTICS

Even when probabilities cannot be calculated with confidence—which, as shown earlier, is the case for most private-market investments—the question of how much to allocate remains. This is where the Kelly Criterion becomes useful, not as a precise formula, but as a heuristic.

## What is the Kelly Criterion?

The Kelly Criterion was developed by John Kelly [Kelly, 1956][29] as a formula for sizing bets when the probabilities and payoffs are known. The formula is:

$$f^* = \frac{p \times b - q}{b}$$

Where:

- $f^*$ = fraction of capital to bet
- ( p ) = probability of winning
- ( q ) = probability of losing (1 − p)
- ( b ) = net odds received on the bet

For example, if an investment has a 60% chance of doubling your money ($b = 1$) and a 40% chance of losing everything, the Kelly formula gives:

$$f^* = \frac{0.6 \times 1 - 0.4}{1} = 0.2$$ or 20%.

But for most family-office investments—private equity, venture capital, direct deals—the probabilities are not known. The outcomes are not bounded. The distribution is not stable. So the formula cannot be applied directly.

## Kelly as a heuristic, not a formula

When probabilities cannot be calculated, some practitioners treat the Kelly Criterion not as a formula to be applied directly, but as a heuristic—a rule of thumb for thinking about sizing. This is the approach taken here.

Under that approach, the sizing question can be framed with four steps:

- Estimate the range of possible outcomes—not precise probabilities, but a rough sense of what could happen.
- Assume the probability of success is lower than you think—the heuristic equivalent of conservatism.
- Size the exposure so that a worst-case outcome does not permanently impair the family—the survival constraint.
- Use half-Kelly as a default—on the basis that it is more robust to estimation errors.

## The half-Kelly heuristic

Half-Kelly means betting half of what the formula would suggest if the probabilities were known. In practice, when probabilities are unknown, half-Kelly becomes a simple rule:

Never allocate more than 10–20% of total capital to any single risk-pole investment.

Why? Because the Kelly formula, even with optimistic assumptions, rarely suggests more than 20–40%. Half of that is 10–20%. This is not a

precise calculation. It is a heuristic that protects against the one-path problem.

Why half-Kelly matters for family offices:

| Feature | Full Kelly | Half-Kelly |
|---|---|---|
| Growth rate | Maximum | Slightly lower |
| Drawdowns | Large (30–50%) | Significantly smaller |
| Recovery time | Longer | Shorter |
| Survival | Less certain in adverse sequences | More certain |
| Robustness to estimation error | Low—errors in p or b change the outcome significantly | High—the buffer absorbs estimation errors |

A family office cannot raise new capital from outside investors the way a hedge fund can. It has one balance sheet, one sequence, one future. Half-Kelly is a response to this. It trades some growth for safety. It preserves the ability to remain in the game.

### The link to the barbell

The barbell strategy is a practical application of the half-Kelly heuristic:

- The safe pole (70–80% of capital) ensures survival—this is the equivalent of not betting.
- The risk pole (10–20% of capital) provides upside—this is the equivalent of the half-Kelly bet.

The sizing of each pole reflects the same principle: bet enough to grow, but not so much that a bad sequence destroys the family. The barbell does not pretend to know the probabilities of the risk pole. It simply sizes the exposure so that the family can survive the sequence of outcomes, whatever they may be.

## 2.7 SUMMARY: WHAT HEURISTICS ACHIEVE

Heuristics achieve the following:

1. They preserve optionality—they keep options open when the future is uncertain.
2. They prevent ruin—they prioritise survival over optimisation.
3. They match the environment—they are tailored to the structure of the problem, not the complexity of the model.
4. They reduce cognitive load—they free attention for what matters most.

The question complexity theory raises is not "How do we predict better?" but rather "How do we make decisions when prediction is limited?" [Gigerenzer & Goldstein, 1996] [Gigerenzer, 2008]. Heuristics are the answer to that question.

## 2.8 THE BRIDGE TO THE BARBELL

The progression leads us here:

a) Complexity tells us why the future is hard to calculate.
b) Uncertainty tells us why probabilities cannot always be assigned.
c) Heuristics tell us how to decide when neither calculation nor probability is enough.

The barbell strategy is a heuristic—a simple rule that matches the structure of the problem family offices face.

# PART 3: THE BARBELL AND SUPPORTING HEURISTICS

## LEVEL 1: THE BARBELL — THE SHAPE, THE OVERARCHING HEURISTIC

The barbell is a shape, not a formula. Imagine an actual barbell—the kind used in weightlifting. It has heavy weights on both ends and a thin, empty bar in the middle. That is the visual the strategy borrows: most of the capital sits at one heavy end; a smaller amount sits at the other; almost nothing sits in the middle.

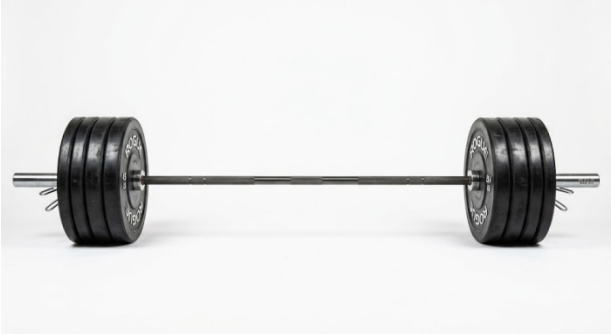

The heavy safe end holds what is intended to still be standing after a hard run: cash, short-dated high-quality bonds, and—if it already produces cash—an operating business the family knows how to run. The other end, smaller but still significant, holds things that can grow a great deal but are volatile: direct companies, venture funds, private equity, permanent capital vehicles. The middle—the thin bar—holds almost nothing. That means avoiding the moderate middle that feels balanced, charges ongoing fees, moves with volatile assets (e.g. unit trusts, balanced funds, funds of hedge funds), and still cannot be sold when cash is needed [Taleb, 2012]. A balanced unit trust charging 1.5% a year is one example: the fee is drawn whether the fund rises or falls, and the units cannot always be sold quickly without penalty.

### Table 1: the barbell ends

| Safe end | Risky end |
|---|---|
| Cash and high-quality bonds [Dew, 2024*][30] | Private equity and venture capital [Dew, 2024*] |
| Blue-chip dividend stocks [Dew, 2024*] | Startup equity [Dew, 2024*] |
| Investment-grade fixed income [Dew, 2024*][30] | Private credit [Dew, 2024*] |
| Mortgage funds [Dew, 2024*] | Development real estate [Dew, 2024*] |

The middle—moderate stocks and bonds—receives less emphasis in this approach [Ruggie, 2026*][31]. The term comes from weightlifting equipment, with weights on both ends of the bar [Wholesale Investor, 2023*][32].

The shape fails if the "safe" end is leverage dressed up as safety, or if the risky end is allowed to become the bulk of the book. The point is survival first, then the chance of a large positive surprise, while avoiding the busy middle that looks adult and still leaves the office trapped.

### Noting the limits of the safe end

The safe end is not assumed to survive every hard run. Cash can be eroded by inflation. Short-dated bonds carry default risk. An operating business can lose customers.

What the safe end offers is not a guarantee. It is access. Its purpose is to remain reachable—to be the part of the book the office can still use when the rest cannot be sold.

That access rests on four properties:

- It can be withdrawn quickly. Cash and short-dated bonds settle in days, not months.
- Its value is known. You can see what it is worth without waiting for a buyer's opinion.

- It does not need a buyer to agree. Government bonds have a market; a private company does not.
- It is not locked. No lock-up period, no gate, no waiting for an exit.

Because of these, the safe end can meet obligations—capital calls, taxes, family needs—without forcing a sale elsewhere at the wrong moment.

That is the job of the safe end: to remain reachable. Not to be untouched by a hard run, but to be available when the office needs to act.

Diversifying the safe end across countries reduces concentration, but it does not isolate the office from the system. Markets are intertwined. A crisis in one country can spread to others through what the literature calls wake-up call contagion—investors reassess countries with similar vulnerabilities and pull capital from them, even where no direct financial link exists [Goldstein, 1998][33] [Ahnert & Bertsch, 2022][34]. The 1997 Asian crisis began in Thailand and spread across the region on this basis. The 1994 Mexican peso crisis hit Brazil hard, even though Brazil's fundamentals had been strengthening [Dasgupta, Leon-Gonzalez & Shortland, 2011][35]. Diversification helps. It does not make the safe end immune.

## An example to think with

Imagine a family office with R100 million in total capital.

### *The safe end (R80 million):*

- R40 million in short-dated high-quality bonds
- R30 million in cash and money-market instruments
- R10 million in a simple, cash-generating operating businesses the family have run for years—a logistics company, a property rental portfolio, or a distribution business they understand intimately

### *The risky end (R15 million):*

- R5 million in venture capital funds
- R5 million in direct investments across a few early-stage companies
- R5 million in private-equity co-investment opportunities

### *The middle (R5 million):*

- A small allocation to balanced mutual funds that charge ongoing fees, move with the broader market, and cannot be sold quickly without penalty

### *Why this is a barbell:*

- The R80 million safe end is sized to absorb a shock. If the entire risky end were lost, the family would still hold R80 million. On that assumption, it could continue to operate, meet obligations and make future commitments. This is not a guarantee: the safe end is not immune to shock. It can be reached by inflation, sovereign default or a currency collapse. It is sized to give the office room to act—not to stand untouched.
- The R15 million risky end is designed to provide exposure to asymmetric upside. One venture investment could return 5× or 10×. But if all three fail, the loss is intended to be painful without being fatal. The design cannot guarantee that outcome; it can only size the exposure so that failure does not threaten the rest of the book.
- The R5 million middle is where the design sees the least value. It offers less safety than the safe end and less upside than the risky

end. It tends to charge fees, move with volatile assets in downturns, and resist quick liquidation when cash is needed. The barbell keeps this space thin for that reason.

*What this is not:*

- It is not a 50/50 split. The asymmetry is deliberate.
- The safe end is not "risk-free." Government bonds carry interest-rate risk. The operating business faces competitive pressure. But these risks are bounded and understood.
- The risky end is not gambling. Each allocation is sized so that failure does not threaten the family's survival.
- The exact percentages are a family decision. Some may choose 70/30. Others may choose 90/10. The direction—most capital safe, a smaller portion risky, almost nothing in the middle—is what defines the shape.

## Why the barbell fits complexity

Complex systems punish concentration and favour the ability to try again. Small events can have large consequences, and large efforts can produce almost nothing. Path dependence means early choices lock later ones. Adaptation means the players change while capital is still committed. Together, these mean a single theme can run through the whole book: what starts in one position can reach others, and the total effect is not proportional to the first loss. The barbell answers by keeping most resources outside the spread of any single theme, and by limiting the size of any single jump so the office remains able to take further opportunities.

# LEVEL 2: SUPPORTING HEURISTICS — THE PRACTICAL HABITS THAT KEEP THE BARBELL HONEST

The barbell was set out above. It is a shape: most capital at the safe end, a smaller portion at the risky end, almost nothing in the middle. That shape answers the problem set out in this paper—that prediction is limited and the office lives one path.

A barbell can drift. The safe end can quietly become leverage dressed as safety. The risky end can grow until it is the bulk of the book. The middle can fill up with products that feel moderate and cost money to hold. Nothing about the shape prevents this.

What follows are habits that guard it. They are drawn from the features of complex systems—limited prediction, single-path risk, adaptation and the need for another attempt. They are offered as a perspective for thinking in the midst of complexity, not as a tested set, and no claim is made that they work together. The office may use one, several, or all of them, as it sees fit.

Each heuristic stands on its own. Some are drawn from published work, and those are cited. Others are the paper's own framing, offered as a way of thinking rather than a finding.

### *1) Respect the single path (redundancy)*

The office lives one irreversible sequence, not the average of many offices. After any plausible shock to the risky sleeve, the larger part of the book must still function and remain able to make decisions. Survival capital is the platform, not decoration.

***In practice: Redundancy***—extra cash, multiple banking relationships, backup operational plans. The cost of redundancy is lower than the cost of being caught without resources [Taleb, 2012].

### *2) Size the risky sleeve so the house remains (half-Kelly)*

The risky end exists to capture large positive surprises. It must be small enough that its total disappearance does not end the office's ability to act. Follow-on needs are counted before the first cheque.

***In practice:*** Half-Kelly—never allocate more than roughly 10–20% of total capital to any single risk-pole commitment when probabilities are uncertain. This protects the one path that actually exists [Kelly, 1956] [Peters & Gell-Mann, 2016].

### *3) Stagger the doors*

Spread commitments across different years and different types of investments, so that liquidity needs and outcomes do not all arrive at the same moment. Place tickets consistently over time rather than concentrating them.

### *4) Prefer what can be understood and sold*

A simple operating business with customers is closer to a financed exit than a complex story only a specialist can value. Listed stocks are a sale button, not a floor. Permanent capital vehicles and long-dated funds are built to stay invested; they are not cash this year.

***In practice:*** Before committing, ask questions—when can this turn into cash the office can use; who has to buy it for that to happen; what else in the book dies if this theme dies; does saying yes today force a larger draw from the safe end later. If the answers are unclear, the ticket is harder to justify [Gigerenzer, 2008] [Simon, 1955]. Less liquid assets typically require more careful sizing [Dew, 2024*].

### *5) Match the book to the generation that must live it*

The Chosen One may tolerate more risk in familiar fields. Later decision-makers may need more liquidity and clearer rules. The barbell can be re-weighted, but the need for a standing house does not disappear.

***In practice:*** Different generations have different time horizons. Shorter horizons emphasise liquidity and preservation; longer horizons can support greater exposure to illiquid, higher-risk positions [Coyle, 2025*] [Dew, 2024*].

### *6) Treat local knowledge as a filter, not a forecast*

Knowing an industry helps the office recognise a simple firm faster and avoid obvious nonsense. It does not turn uncertainty into risk. Local knowledge is useful; it is not a crystal ball [Hayek, 1945][36].

### *7) Watch the middle that pretends to be safe (via negativa)*

Products that feel moderate, charge for activity, fall when risk assets fall, and still cannot be sold quickly are often the most expensive place to sit. Complexity makes the moderate middle less stable than it looks. The barbell deliberately leaves that space thin.

***In practice:*** Via negativa—knowing what to avoid is often more important than knowing what to pursue. By removing fragility, the office becomes more robust [Taleb, 2012].

## Supporting: frameworks for capital

The following illustration provides broader frameworks for thinking about capital allocation. They are not tied to a single heuristic but underpin the barbell's logic.

## Heuristic: consider the purpose of capital

One way to think about capital is to separate it by purpose [Coyle, 2025*]:

| Capital type | Purpose | Examples |
|---|---|---|
| Survival capital | Meet obligations, provide liquidity | Cash, reliable high-quality fixed income |
| Exploration capital | Pursue asymmetric opportunities | Private equity, venture capital, intellectual property |

This separation is about intent—not about declaring one type good and another bad [Coyle, 2025*]. It provides a mental framework for the barbell's two ends.

# THE ARGUMENT FOR SEEING THEM TOGETHER

The heuristics presented in this section are not a checklist to be applied mechanically. They are interrelated perspectives that reinforce the protection of the family office.

1) They share a common foundation.
   All the rules and heuristics arise from the same recognition: prediction is limited, complexity is the setting, and the family office lives one path through time. Whether the rule is "respect the single path" or the heuristic is "half-Kelly sizing," the underlying logic is identical—survive first, then grow.
2) They address different levels of decision-making.
   The rules operate at the level of specific commitments—how to size a ticket, when to say no, how to stagger doors. The heuristics operate at the level of broader framing—how to think about capital, liquidity and time.
3) They protect against different kinds of failure.
   Different rules and heuristics guard against different failure modes: over-concentration, illiquidity, generational mismatch, false safety, poor governance and inadequate survival capital.

# WHAT THIS PAPER DOES NOT CLAIM

It does not name a winning mix of assets.
It does not say which jurisdictions to invest in.
It does not say how to diversify.
It does not treat listed stocks as a guarantee.
It does not say yearly tickets will exit the following year.
It does not give mechanics of how to structure a family office. Even though some family offices establish multiple companies and trusts to protect their wealth by ring-fencing [Morgan Lewis, 2026*][37] [Appleby, 2020*][38], the paper applies to overall family-office structures—whether a single structure or multiple for ring-fencing.

It only says the future is closed, complexity is the setting, and a barbell with clear rules is one coherent way for a family office to stay able to act.

# CONCLUSION: A FRAMEWORK, NOT A PRESCRIPTION

This guide has offered a perspective—not a set of recommendations.

The framework is not an argument for or against any particular investment approach. It is an invitation to think about capital allocation in a way that acknowledges uncertainty [Knight, 1921].

The final question is not "Which strategy is right?" but rather: given your family's purpose, time horizon and circumstances, what perspective best guides your decisions?

# APPENDIX: THE STATUS OF THE HEURISTICS

The heuristics in this framework are not presented as empirical evidence that a particular allocation will produce a particular outcome. They are propositions for decision-making under uncertainty.

This distinction matters because evidence of investment effectiveness is itself subject to the problem of induction. To observe that an allocation was followed by an outcome is one thing. To infer that the allocation caused the outcome—or that it should be preferred by another family office—is a further step. The movement from observed instances to claims about unobserved ones is the classical problem of induction [Hume, 1748][39].

The problem is sharper for family offices because they differ. Purposes, time horizons, liquidity needs, generational stage and capabilities vary. An outcome observed in one office does not, by itself, establish what will occur in another. Nor is the available evidence usually a complete record: the researcher may not have observed every decision, alternative or contextual change that shaped the result. What is available is an account of a process, not the process itself. The proposition that a strategy caused an outcome therefore involves interpretation beyond the observation.

This is not an argument against evidence. It is an argument against treating evidence as if it eliminates uncertainty. Evidence can establish what was observed under particular conditions. It does not establish a universally optimal allocation.

This is central to the status of the heuristics here. The half-Kelly heuristic is not a claim that 10–20% is optimal, or that 5% or 30% would be worse. It is a deliberately simple default: no single commitment should be large enough that its failure threatens the office's ability to act. The precise figure may vary by family and circumstance. The heuristic places a design constraint on the book—if one commitment is to consume a large share of the risky end, that allocation requires justification.

The same applies to the other heuristics. They are design propositions for decisions under uncertainty,

not universal causal laws. The purpose of the framework is not to replace uncertainty with certainty. It is to offer simple principles for deciding within it while preserving the family's ability to act.

# CONTACT INFORMATION

**Startup Picnic Venture Studio (SPVS)**

**Johannesburg, South Africa, 2090**

**Email: info@startuppicnic.co.za**

**For more information**

**Websites:**

**www.startuppicnic.co.za**
**www.tiisetsomaloma.co.za**
**www.startuppicnicacademy.co.za**

# FOOTNOTE REFERENCES

[1] [Deloitte, 2023*] Deloitte. (2023). Family offices: Preserve, protect and grow your legacy.

[2] [PwC, 2021*] PwC. (2021). What type of family office is right for you? and Family offices: Fundamentals of family ownership vehicles.

[3] [Deloitte, 2026*] Deloitte. (2026). Family office services and operating models.

[4] [Family Office Exchange, 2026*] Family Office Exchange. (2026). What is a family office? and SFO vs. MFO.

[5] [J.P. Morgan, 2026*] J.P. Morgan. (2026). Multi-family vs. single family office.

[6] [U.S. Bank, 2026*] U.S. Bank. (2026). Is a virtual family office right for you?

[7] [Taleb, 2007] Taleb, N. N. (2007). The black swan: The impact of the highly improbable. Random House.

[8] [Arthur, 1999] Arthur, W. B. (1999). Complexity and the economy. Science, 284(5411), 107–109.

[9] [Mitchell, 2009] Mitchell, M. (2009). Complexity: A guided tour. Oxford University Press.

[10] [Hsu & West, 2025*] Hsu, J., & West, J. (2025). Complexity and investors: Navigating an intricate system. [Industry / working paper].

[11] [Holland, 1992] Holland, J. H. (1992). Adaptation in natural and artificial systems. MIT Press.

[12] [Coyle, 2025*] Coyle, J. (2025). Household-first approach to multigenerational wealth. Libretto (cited in WMI Impact, 2025).

[13] [Cambridge Associates, 2023*] Cambridge Associates. (2023). Family office costs and complexity: Managing the trade-off.

[14] [Dijkstra, 1989] Dijkstra, E. W. (1989). On the cruelty of really teaching computing science. Communications of the ACM, 32(12), 1398–1404.

[15] [Wolfram, 2008*] Wolfram, S. (2008). Contribution to Complexity: 5 Questions.

[16] [Simon, 1962] Simon, H. A. (1962). The architecture of complexity. Proceedings of the American Philosophical Society, 106(6), 467–482.

[17] [Arthur, 1989] Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. The Economic Journal, 99(394), 116–131.

[18] [David, 1985] David, P. A. (1985). Clio and the economics of QWERTY. American Economic Review, 75(2), 332–337.

[19] [Knight, 1921] Knight, F. H. (1921). Risk, uncertainty and profit. Houghton Mifflin.

[20] [Gigerenzer & Goldstein, 1996] Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. Psychological Review, 103(4), 650–669.

[21] [Gigerenzer, 2008] Gigerenzer, G. (2008). Rationality for mortals: How people cope with uncertainty. Oxford University Press.

[22] [Peters, 2019] Peters, O. (2019). The ergodicity problem in economics. Nature Physics, 15, 1216–1221.

[23] [Peters & Gell-Mann, 2016] Peters, O., & Gell-Mann, M. (2016). Evaluating gambles using dynamics. Chaos, 26(2), 023103.

[24] [Gigerenzer, Hoffrage & Goldstein, 2008] Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition. Psychological Review, 115(1), 230–239.

[25] [Tversky & Kahneman, 1974] Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185(4157), 1124–1131.

[26] [Simon, 1990] Simon, H. A. (1990). Invariants of human behavior. Annual Review of Psychology, 41, 1–19.

[27] [Simon, 1955] Simon, H. A. (1955). A behavioral model of rational choice. Quarterly Journal of Economics, 69(1), 99–118.

[28] [Taleb, 2012] Taleb, N. N. (2012). Antifragile: Things that gain from disorder. Random House.

[29] [Kelly, 1956] Kelly, J. L. (1956). A new interpretation of information rate. Bell System Technical Journal, 35(4), 917–926.

[30] [Dew, 2024*] Dew, J. (2024). Billionaire wealth strategies. [Billionaire Investment Allocation Model].

[31] [Ruggie, 2026*] Ruggie, T. (2026). The barbell era: How ultra-wealthy investors are positioning for what comes next. Destiny Family Office, April 2026.

[32] [Wholesale Investor, 2023*] Wholesale Investor. (2023). [Industry commentary on barbell terminology].

[33] [Goldstein, 1998] Goldstein, M. (1998). The Asian financial crisis: Causes, cures, and systemic implications. Institute for International Economics.

[34] [Ahnert & Bertsch, 2022] Ahnert, T., & Bertsch, C. (2022). A wake-up call theory of contagion. Review of Finance, 26(4), 829–854.

[35] [Dasgupta, Leon-Gonzalez & Shortland, 2011] Dasgupta, A., Leon-Gonzalez, R., & Shortland, A. (2011). Regionality revisited: An examination of the direction of spread of currency crises. Journal of International Money and Finance, 30(5), 831–848.

[36] [Hayek, 1945] Hayek, F. A. (1945). The use of knowledge in society. American Economic Review, 35(4), 519–530.

[37] [Morgan Lewis, 2026*] Morgan Lewis. (2026). Building the foundation: Legal architecture for the modern family office.

[38] [Appleby, 2020*] Appleby. (2020). We are family and the vital role of the family office. Business Brief.

[39] [Hume, 1748] Hume, D. (1748). An enquiry concerning human understanding.