Ethical Dilemmas in AI and Financial Services

First HSBC-GCD Conference for Credit and Financial Risk in Krakow

Dr. Philippe J.S. De Brouwer

prepared for: HSBC

September 9, 2026

Philippe De Brouwer

Academic and Professional Profile Dr. Philippe De Brouwer holds a background in theoretical physics and earned a second Master's degree in Business Engineering while working full time. His academic distinction includes solving the "fallacy of large numbers" puzzle, originally posed by P.A. Samuelson and unresolved for 38 years. In his Ph.D., he challenged the foundational assumptions of Harry Markowitz's Nobel Prize-winning Mean Variance Theory, developing the "Maslowian Portfolio Theory" and fundamentally advancing the field of investment advice.

Career Highlights Early in his career, De Brouwer transitioned from insurance to banking and from IT to asset management. At Fortis (now BNP Paribas), he played a key role in launching one of Belgium's first capital-guaranteed funds and was promoted to director in 2000. In 2002, he joined KBC, where he orchestrated the merger of four companies and became CEO of the consolidated entity in 2005. Under his leadership, the company climbed from 11th to 5th in the market, even amid heightened competition.

Following the financial crisis, De Brouwer led the creation of a new investment management company for KBC in Ireland, which quickly grew to manage around 1000 investment funds and approximately €32 billion in assets. In 2012, he broadened his expertise into financial risk management, focusing on statistics, analytics, data, and numerical methods.

Recent Roles and Publications De Brouwer served as Head of Analytics Development for the Royal Bank of Scotland Group and currently holds the position of SVP at HSBC in Kraków. There, he oversees the Model Review Centre of Excellence and serves as Global Head of Professional Practices for Model Risk Management. He also coordinates collaboration with universities.

He has authored numerous scientific papers and books, with his latest being "The Big R-Book: From Data Science to Big Data and Learning Machines." His work continues to influence both academic and practical approaches to finance, risk, and data science.

Wikipedia
Website

Contact
LinkedIn
e-mail
mobile: +48 790 715 002

Contents

1 Defining Ethics

A Formal Definition – Cambridge Dictionary

Ethics

the study of what is morally right and wrong, or a set of beliefs about what is morally right and wrong.

Britannica: ethics, also called moral philosophy, the discipline concerned with what is morally good and bad and morally right and wrong. The term is also applied to any system or theory of moral values or principles.

How should we live? Shall we aim at happiness or at knowledge, virtue, or the creation of beautiful objects? If we choose happiness, will it be our own or the happiness of all? And what of the more particular questions that face us: is it right to be dishonest in a good cause? Can we justify living in opulence while elsewhere in the world people are starving? Is going to war justified in cases where it is likely that innocent people will be killed? Is it wrong to clone a human being or to destroy human embryos in medical research? What are our obligations, if any, to the generations of humans who will come after us and to the nonhuman animals with whom we share the planet?

Ethics deals with such questions at all levels. Its subject consists of the fundamental issues of practical decision making, and its major concerns include the nature of ultimate value and the standards by which human actions can be judged right or wrong.

The terms ethics and morality are closely related. It is now common to refer to ethical judgments or to ethical principles where it once would have been more accurate to speak of moral judgments or moral principles. These applications are an extension of the meaning of ethics. In earlier usage, the term referred not

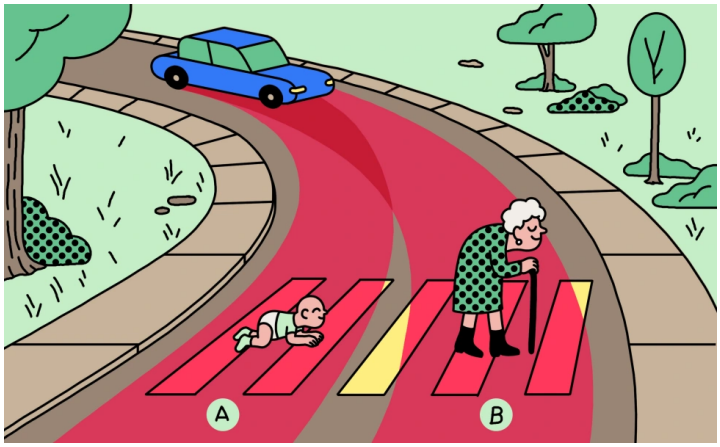

Figure 1: *You can only correct the course of the car a little, so you can reduce the deaths from two to one. Do you choose track A or B? — source: MIT Technology Review — Hao, 2018*

to morality itself but to the field of study, or branch of inquiry, that has morality as its subject matter. In this sense, ethics is equivalent to moral philosophy.

Although ethics has always been viewed as a branch of philosophy, its all-embracing practical nature links it with many other areas of study, including anthropology, biology, economics, history, politics, sociology, and theology. Yet, ethics remains distinct from such disciplines because it is not a matter of factual knowledge in the way that the sciences and other branches of inquiry are. Rather, it has to do with determining the nature of normative theories and applying these sets of principles to practical moral problems.

This article, then, will deal with ethics as a field of philosophy, especially as it has developed in the West. For coverage of religious conceptions of ethics and the ethical systems associated with world religions, see Buddhism; Christianity; Confucianism; Hinduism; Jainism; Judaism; Sikhism.

Ethical dilemmas are more important than ever. With the next figures we explore a few variations of the trolley problem. The Trolley Problem is no longer just a thought experiment for philosophy students; it is a technical requirement for autonomous systems. Engineers must decide how a car should value a child's life versus an elderly person's life in a situation where a collision is unavoidable.

Dilemma 1: Kill the Baby or Old Person?

It appears that the choices people make is somehow determined by culture. In Western culture a majority of people will choose to save the baby, in West Asian cutltures typically the older person is saved.

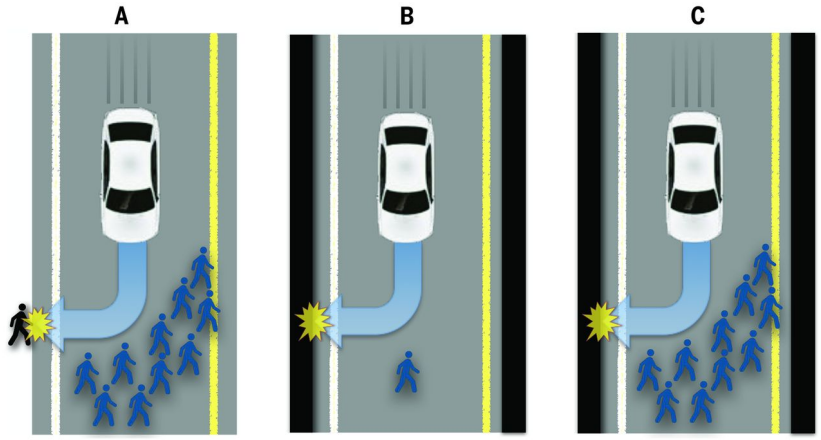

Figure 2: *source: Science — Bonnefon, Shariff, and Rahwan, 2016*
Should we veer and kill the person who is on the pavement? What if there is only one person on the street and there is a wall? How many people to change behaviour?

More Dilemmas

Is it ethical to remove a tree that houses a squirrel to build your house?

Figure 3: *Is removing a tree that houses a squirrel ethically correct when building your home? What if an advanced society wants to make Earth their home?*

Most people will find that removing the tree that houses the squirrel is morally acceptable. If the squirrel would be another human, typically (in modern times) we would not accept this. However, we and the squirrel share a common ancestor that lived ca. 180 million years ago.

We are only separated by 180 million years, and we do not see the squirrel as a valid discussion partner for the location of our house. Now, imagine an alien society that has a 180 million years earlier start of their civilisation than us. They will probably not recognise us as intelligent partners, and just like the squirrel we will be defenceless against their technology if they want to take our planet.

Suddenly the same problem feels very different.

What is ethically acceptable depends on your viewpoint

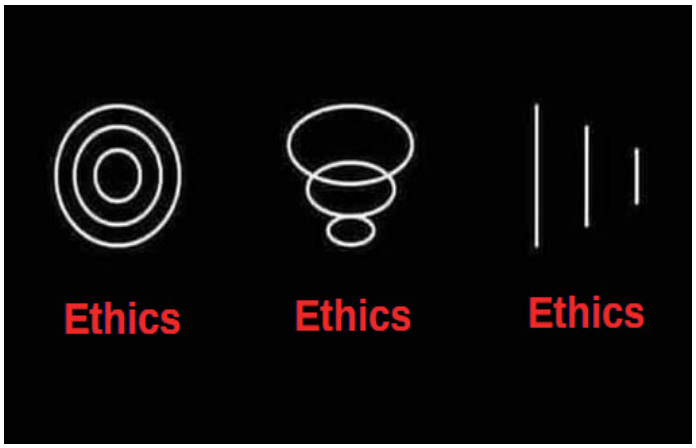

Figure 4: *Ethics is complex and it depends on your point of reference and you frame.*

2 Artificial Intelligence

A Definition of AI

Artificial Intelligence

A cross-disciplinary approach to understanding, modelling, and replicating intelligence and cognitive processes by invoking various computational, mathematical, logical, mechanical, and even biological principles and devices. —Brundage, 2015

or

AI in Britannica.com

the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings

or

AI in Wikipedia

Artificial intelligence (AI) is intelligence demonstrated by machines, as opposed to the natural intelligence displayed by animals and humans.

First Some Simple Questions

- What is exactly Intelligence and how to measure it?

  the ability to acquire and apply knowledge and skills (Oxford Dictionary)

- What is exactly consciousness?

  the state of being aware of and responsive to one's surroundings - Oxfor Dictionary

- What is Self-consciousness? Sentience? Life?

Issues with Morality and AI

- What if AI becomes more intelligent than people? Does it get moral status?
- What about the singularity and benevolence of machines?
- What about enhanced humans (e.g. math implant in brain?)
- What when we can "upload" a human brain? Who owns it? Can we switch it off?
- Does a conscious and sentient AI deserve rights?
- Who are we to decide this?
- Can one delegate agency to an AI? Who has the moral agency?
- Should we manipulate data for "good"? (and counter-act bias?)

3 Example: Insurance Onboarding

Car Insurance Claims Data
The data contains

- driving related:
  - average daily travel
  - professional or private use
  - living in urban or city area
- car related:
  - type of car (eg. minivan, sports car, sedan, etc.)
  - age of the car
- behaviour on our books:
  - time with us
  - old claims
- driver related:
  - has driver's licence been revoked
  - penalty points on driver's licence
  - years on job (work related)

. . . but also contains

- related to gender and family:
  - number of kids at home
  - single parent or not
  - marital status
  - gender
- education and profession
  - education
  - occupation
- wealth related:
  - value of the car
  - value of the home

We call these the **"protected features"**

Protected features are variables that, for ethical or legal reasons, should not influence a decision-making model. In insurance, using factors like gender or marital status can lead to discriminatory pricing, even if those factors are statistically predictive of risk.

The cost of being politically correct?

Area Under the Curve (AUC) is a foundational statistical metric used to evaluate the performance of binary classification models, particularly the Receiver Operating Characteristic (ROC) curve in risk analysis. Within the context of insurance underwriting, AUC measures a model's discriminative power—its structural capability to correctly rank-order policyholders so that individuals who will eventually file an accident claim receive higher risk scores than those who will remain claim-free.

- **Discriminative Power:** AUC quantifies the probability that a randomly selected high-risk customer is assigned a higher risk score than a randomly

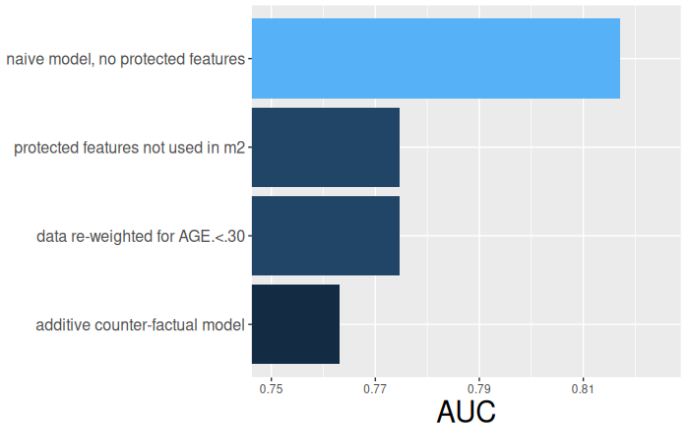

Figure 5: *The discriminating power of different models. The higher the AUC, the stronger the model, the lower the losses on the portfolio, the higher profit (or lower prices) can be.*

chosen low-risk driver. An AUC score of 0.5 denotes performance no better than random guessing, whereas a score of 1.0 indicates a theoretical, flawless predictor.

- **Portfolio Efficiency:** Stronger models with higher AUC scores allow insurers to fine-tune pricing tiers, minimize cross-subsidization, reduce portfolio-wide losses, and offer more competitive pricing to safe drivers.
- **The Ranking Limitation:** AUC evaluates *relative ordering* rather than absolute risk probability or financial magnitude. It confirms whether high-risk drivers are positioned above low-risk drivers, but it remains entirely blind to the actual monetary value of potential claims.

Because AUC is strictly a rank-order metric, it cannot be directly converted into a precise monetary figure or a specific financial impact. When compliance frameworks or ethical mandates require data scientists to strip out protected features—such as gender, marital status, or socioeconomic indicators—the model's AUC inevitably drops, representing a loss in statistical precision.

However, translating this drop in AUC into an exact dollar amount or premium hike for an individual policyholder is fundamentally impossible through the metric alone. A reduction in AUC from 0.85 to 0.80 indicates degraded portfolio discrimination, but it does not dictate how pricing tables must shift.

Quantifying the financial cost of eliminating bias requires mapping the altered risk distributions against commercial underwriting rules, loss cost projections, and regulatory pricing constraints. While a weaker model inherently introduces portfolio inefficiencies, the exact financial burden experienced by specific consumer groups remains concealed behind complex actuarial formulas rather than visible within the AUC score itself.

How to avoid bias? Two alternatives

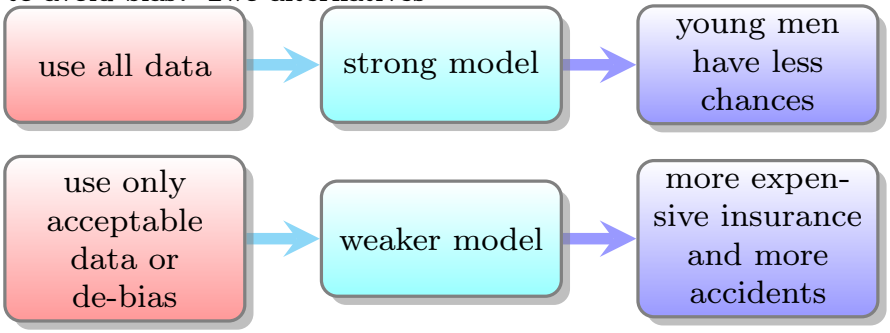

4 Ethics in Data Science

Some Issues

- user consent
- privacy, data protection
- intellectual property rights, data ownership
- safekeeping: avoiding data theft, cyber security
- trust, surveillance, free will
- competition
- discrimination:
  - bias
  - fairness
  - accountability
  - transparency

As AI systems gain agency, we face the "Alignment Problem" –ensuring that machine goals remain compatible with human values. This includes determining who is legally and morally responsible when an autonomous system makes a harmful decision.

5 Conclusions

Visual Bias 1

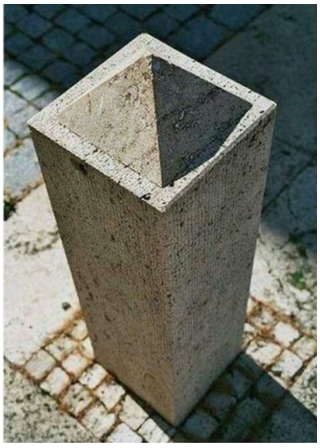

Figure 6: *Visual biases as clear: is the pole hollow or is there a pyramid on top of it?*

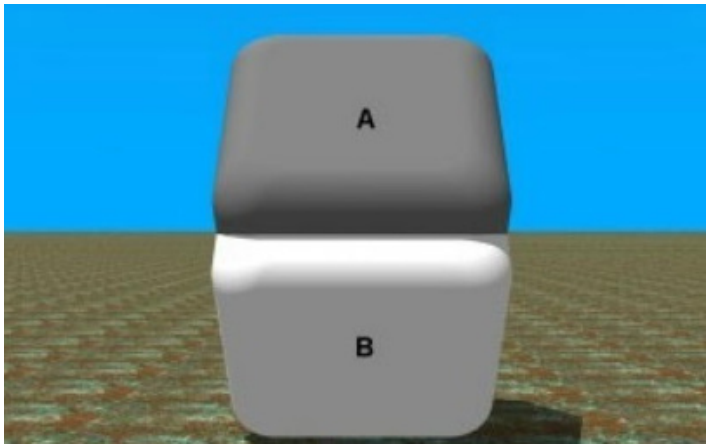

Figure 7: *Visual biases cannot be avoided. We know that A and B have the same shade of grey, but we cannot see it correctly.*

Visual Bias 2

Fairness is a matter of perspective

We do have an innate sense for what is fair and what not, however that sense is heavily biased towards ourselves. For example, we don't see it as a moral issue to build a house and destroy the home of a squirrel in the process. The squirrel, however, is only separated from us by 180 million years of separate evolution. From a society that is a million years ahead of it might seem obvious to simply make earth their home and get rid of us.

If that seems unfair, turn it around. Would it be fair to colonise a planet inhabited by squirrel-like mammals and make it our home?

Figure 8: *Fairness is a matter of perspective*

Bias is learnt

Figure 9: *Bias is assuming a solution based on what you have learnt before. It can be useful, but usually is limiting.*

Conclusions

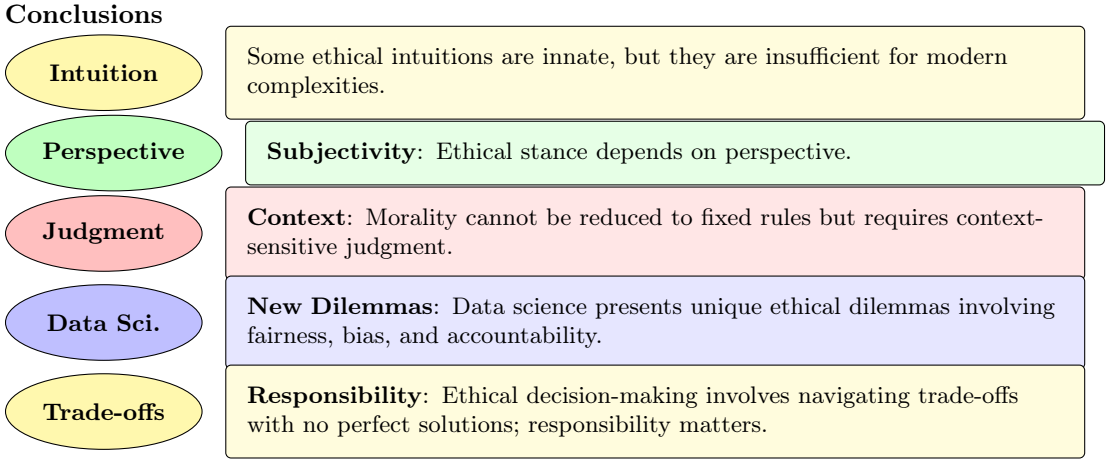

References

Bonnefon, Jean-François, Azim Shariff, and Iyad Rahwan (2016). 'The social dilemma of autonomous vehicles'. In: *Science* 352.6293, pp. 1573–1576.

Brundage, Miles (2015). 'Taking superintelligence seriously: Superintelligence: Paths, dangers, strategies by Nick Bostrom (Oxford University Press, 2014)'. In: *Futures* 72, pp. 32–35.

Hao, Karen (2018). 'Should a self-driving car kill the baby or the grandma? Depends on where you're from'. In: *URl:https://www.technologyreview.com/s/612341/a-global-ethics-study-aims-to-help-ai-solve-the-self-drivingtrolley-problem/*.