Chapter 2: How do people make ethical decisions, and how should they?

Abstract: Before looking at how people should make ethical decisions, it is useful to start by looking at how people normally do make ethical decisions. This has been the subject of much research in behavioural economics, psychology, social psychology and sociology. This research suggests that many ethical decisions are made intuitively. These intuitions are very helpful as they allow 'fast' thinking, which means people can function without being paralysed by the need to think everything through. But intuitions are also subject to predictable errors. These errors include priming effects, anchoring bias, a disposition to follow narrative over logic, and bystander effects. Recognising how these biases impact ethical thinking helps clarify how to avoid them.

Case Vignette: Budget estimation

Dominique is designing a research study that uses animal models to test the efficacy of a particular treatment. The study design is similar to a number of previously published studies. The most cited of these studies involved n=172 rabbits and found positive outcomes in 90% of cases in the experimental sample compared to 60% in the control group. Other studies had similar designs and outcomes. A critical significance value (p-value) of 0.05 and a power of 80% are accepted as the standard.

Dominique is working on the budget for the new study and needs to estimate how much money to allocate to animal testing. The actual number of animals to be used can be determined later – for budgetary purposes only an estimation of the number of animals, and hence the costs, is needed. What estimate will Dominique make? Enough money for:

a. 60 rabbits?
b. 110 rabbits?
c. 160 rabbits?
d. 210 rabbits?

Keep a note of your answer – we will return to this later.

Introduction

In this chapter, we will focus on the thinking process through which people come to ethical decisions. This topic can itself be approached through (at least) two questions: the first is to ask "How *should* people make ethical decisions?" (this is referred to as *normative* ethics). The second is to ask, "How *do* people make ethical decisions?" (this is referred to as *descriptive* ethics). One of the key ideas in ethics education is that the way people *do* make decisions often does not match the way that ethicists say they *should* make decisions. Therefore, a purpose of ethics education is to teach a way of making ethical decisions, and to allow people to practice that decision-making approach often enough that it becomes something they can easily apply when they need to. But before we address what we should do, it makes sense to think a bit further about some of the problems arising from our more habitual decision making process.

Intuitions and ethical decision making

Quickly answer the following question:

> *A bat and a ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost?*

Think back on what happened in your head when you saw the question. For many people, their first impulse is to intuitively answer '10 cents'. Some will then recognise that this answer is wrong (they may recognise that the question is framed like a trick question, which might cause them to question their initial answer). They will then work out the answer by solving the two simultaneous equations embedded in the question. But a lot of people won't give their intuitive answer a second thought and will simply answer '10 cents'.

This ‘trick question’ is actually one of a series of questions which make up part of a psychological test called the Cognitive Reflection Test (CRT), developed at MIT. The CRT is used to assess whether or not a person tends to accept their initial intuitive (‘gut’) answer to certain types of questions, or whether they reflect on, and correct, their intuitive answer. The CRT contains three questions designed to give rise to an intuitive answer like the one above. Results using the CRT found that out of a sample of over 3,400 people, 83% of respondents did not reflect on their answer in at least one of the three questions. 33% of respondents didn’t reflect on their answer in any of the three questions. In studies with the CRT in different universities, the group which scored highest on checking their answers were the students in MIT, but, even there, over half of the respondents didn’t reflect on or adapt their intuitive answer (Frederick, 2005: 29) on at least one of the three questions.

This kind of ‘trick’ question works so often because it leverages a particular feature of human thinking, which is referred to as *heuristic* thinking. A heuristic is “a simple procedure that helps find adequate, though often imperfect, answers to difficult questions” (Kahneman, 2011: 98). To use an analogy from computing, in situations in which it is difficult to arrive at a perfect solution to a problem, the term heuristic is applied to a method which yields an approximate solution that is regarded as ‘good enough’ to within a given margin of error. The term is used similarly in relation to cognition: heuristics are ‘shortcuts’ that our brain uses to arrive at a ‘workable’ solution in situations where an accurate solution may take additional time or too much effort. In the case of the bat and ball ‘trick’ question, the question is phrased in such a way as to trigger a cognitive heuristic process whereby our brain takes a complex question which requires manipulating two variables and two equations and replaces it with a simpler question which involves a simple subtraction and which can be answered without conscious effort by the automated functions of our cognitive processes.

A key point of research into such thinking heuristics is that people use them all the time without necessarily being aware that their decision results from a heuristic. The research findings on cognitive heuristics are relevant to ethical decision making since our brain does not only use the 'solve a simpler problem' heuristic when dealing with word-based mathematical problems. It also uses the same heuristics when faced with ethical problems which are often, almost by definition, complex questions. To take one example, one cognitive bias that has been identified by researchers is *extension bias* which refers to a tendency to assume that more is better than less. This has been identified as having a range of impacts on bioethics. For practitioners and study designers it has been seen in a tendency to chose more tests or examinations over fewer when given the choice. For those working in areas such as genetic engineering or in human enhancement it is seen in a tendency to unquestioningly assume that more of a specific feature (intelligence, speed, strength…) is better than less. It has also been identified as providing a flaw in ethical reasoning in that the more arguments we find for a position the stronger the position is seen as being (irrespective of the quality of the arguments) (see Hofmann, 2023).

Indeed, the 'solve a simpler problem' heuristic is only one of the mental shortcuts that may arise in the context of moral decision making. Another is called 'anchoring'. Consider the following two questions (which comes from Kahneman, 2011: 43):

1. *Is the height of the tallest redwood tree more or less than 1,200 feet (366 m)?*
2. *What is your best guess about the height of the tallest redwood?*

The answer to the first question is pretty straightforward even if you don't know much about trees ('probably not taller than 366 m'). In experimental studies, the mean average guess

people give for the second question is 844 feet (257 metres). However, if the first question is changed slightly to '*Is the height of the tallest redwood tree more or less than 180 feet (55m)?',* then the mean average guess people give for the height of the tallest redwood is 282 feet (86 metres). Although the second question is the same in both cases, the framing provided by the first question changes the answer to the second question, with the difference between the two estimates being an incredible 562 feet (171m), (which is, ironically enough actually substantially taller than the tallest redwood tree!).

This anchoring effect can have a significant impact on ethical decision making. For example, in one study, visitors to an environmental park were asked to donate money to protect seabirds from environmental damage caused by oil spills. Three conditions were used in the study. In a neutral condition, participants were simply asked to donate money. In a 'low anchor' condition, they were asked, "Would you be willing to pay 5$..." and then asked to donate. In a 'high anchor' condition, they were asked, "Would you be willing to pay 400$..." and then asked to donate. In the neutral condition, the average suggested donation was 64$. In the low anchor condition, the average suggested donation was 20$. In the high anchor condition, it rose to 143$ (Kahneman, 2011: 124-5). Anchors probably operate by (subconsciously) activating a series of memories and ideas linked to the suggestion, which then affects the thinking process for the second task.

Case Vignette (continued): Budget Estimation

1. Think back on your answer to the vignette that started this chapter. Most people chose either 160 or 210 rabbits. This makes sense in the context of the 'anchoring' heuristic – people intuitively start from 172 and work up or down from there. When you made the decision, what justification did you have for that decision? Remember that heuristics

operate pre-cognitively (i.e. without the person being aware), so the justification you give yourself for your answer may not reflect the actual operation of the heuristic: did your explanation to yourself of your answer include an understanding of heuristics?

2. For this case, try to identify the range of different entities which are affected by Dominique's decision: (these may be humans or non-humans, those near to Dominique and those far away….). To what extent were these different stakeholders considered in your answer to the initial question?
3. For three of the entities most affected, describe in a sentence or two how they would describe the research project *from their perspective*. What emotions might they feel about Dominique's decision?
4. Are there particular statistical methods that could be used to estimate the sample size needed for an experiment or a study? How might a life scientist use their statistical competence to arrive at a better estimation?
5. Based on the data available, a reasonable estimation of the size needed for the experiment is 62 rabbits (31 in each of the control and experimental group). One of the ethical principles applied in bioscience is the principle of 'non-maleficence': the principle of minimising harm. In animal testing this is operationalised as a practice of 'reducing' – i.e. causing harm to as few animals as is necessary for the research. Does the 'anchoring' heuristic in this case result in an ethical or un-ethical outcome from animal testing?
6. What kind of procedure might be put in place in the lab to ensure this kind of decision is made in future in ways that ensure care for all those who need it, in this context?

Anchors can be thought of as one type of what is referred to as a *priming* effect, in which an idea that is brought to mind then affects subsequent decision making. For example, the presence of ideas about money has been found to affect how people subsequently make

decisions. In studies, various methods have been used to ensure people were subconsciously aware of money: these include asking them to make short sentences out of words which were selected to draw their attention towards money, or even by simply leaving a stack of monopoly money in their sightline. The effect of this 'money priming' on their actions was notable. While money-primed people persisted longer at tasks, they also gave, on average, less help to others, and when asked to set up two chairs for a conversation, money-primed people set the chairs further apart than non-money-primed people (Vohs, 2006). The simple presence of the idea of money seems to encourage people to act in more individualistic and selfish ways. Again, this is not something the people in question would even be aware of.

Now have a go at answering the following question, before reading further:

> *Linda graduated with a Master's degree in English literature 15 years ago. Her thesis was on representations of female domestic servants in 19th Century American novels. What is her job today? (Rank in order of likelihood from 1 [most likely] to 5 [least likely])*
>
> a. *Manager of a number of crèches*
> b. *Bank worker*
> c. *Owner of a small, independent bookshop who organises women's literary evenings*
> d. *Bank worker who volunteers with an organisation offering support to women-headed small businesses*
> e. *Head of a cantonal public library board*

While the detail of the story focuses on Linda's own experiences, the probability of her having a certain job also depends on something the story did not focus on: the frequency of those jobs within the labour market. It is probably evident that there are far more jobs working in banks than there are, for example, as head librarian in a cantonal library. Despite this, many people will list *e* as being more likely than *b*. The question gets even more interesting if we look more closely at the jobs *b* and *d*. It is useful to think of these two categories as being sets. As sets, we can immediately recognise that the category of *d* is a

sub-set of *b*, that is, category *b* contains all cases of *d* plus additional cases. Therefore, by definition, *b* is more probable than *d*. Yet when asked about Linda, many people chose the *d* as being more likely than *b*. Why do so many people choose an answer that is easily identified as wrong if one applies some basic mathematics to the question? In this case, the heuristic 'simpler question' that our brain asks when faced with a slightly complex question requiring estimation of unknown parameters in determining probability is 'in this story, what happens next?'. The story of Linda working in a job that focuses literature or on women and their advancement seems like a more compelling and satisfying conclusion to her trajectory than does the story of her (just) working in a bank. Hence the librarian and the volunteer stories intuitively seems more probable than her (just) working in a bank, even if some simple maths would demonstrate otherwise. This example again is loosely based on one used by Kahneman. He explains this as another of the predictable errors that arise from our brain's heuristics, a heuristic that he calls a tendency to 'narrative (story) over numbers'.

The story of 'Linda' also demonstrates another heuristic that affect decision making. In this story, there was information about Linda's prior academic work, but none about the labour market in the country where she lived or the number of jobs available in libraries or banking. When faced with situations (like this) of partial or incomplete evidence, the human brain will often proceed on the basis that it has all the information necessary to make a decision. Kahneman (2011) calls this the 'what you see is all there is (WYSIATI)' heuristic. Applied in the context of ethical decision making, this heuristic could lead someone to come to a judgement based on the information in front of them, rather than to search more widely to identify additional information that may be relevant in the judgement. Think back to the 'budget estimation' case vignette: did you initially consider the impact of your decision on the animals that would be involved in the testing?

Case History: Where do antibodies come from?
The WYSIATI heuristic

Students in Life Science engineering use antibodies in the course of their studies. In EPFL, for example, they are used in courses like 'Integrated labo in Life sciences II' (Bio-204). In this course, for example, students perform a western blot to test whether a protein of interest was produced by cells. Antibodies used for western blots are sourced from companies that produce them (such as, for example, Thermo Fisher Scientific) and come in sterile plastic bottles. There are two ways of making antibodies:

1) Classical by injecting live animals and collecting the blood (usually several bleeds before sacrificing the animal)
2) In cell culture a) by making cell lines from a previously injected mouse or by b) genetic engineering

The Thermo Fisher Scientific website seems to indicate that they primarily use the classical method: "Multiple species are used to generate antibodies that can be used in western blot applications. Most commonly: mouse, rabbit, rat, goat, donkey and chicken… however, most primary research antibodies for western blotting are produced from immunized rabbits (polyclonal, monoclonal, recombinant) or mice (hybridoma derived monoclonals). Some host species provide additional advantages over others, for example, due to their size or immune biology. For example, when comparing mouse or rabbit, rabbits usually are better at tolerating immunizations and have a significantly longer life span than mice. Furthermore, rabbits exhibit a more diverse natural repertoire of antibodies than mice, which makes rabbits a popular host for the generation of polyclonal, monoclonal and rabbit recombinant antibodies" (link).

Two of the students who helped to develop material for this course remember taking the Bio-204 course and working with antibodies. Although the information about the origin of the antibodies was available, it was not something that the students particularly paid attention to. VB said "I do not remember asking about it or feeling weird about it, since by then we had worked for an entire semester (almost 2) in the wet lab, so we were used to handling cells and [handling antibodies] did not feel any different. They [i.e., the antibodies] were also more presented as a tool to perform western blot". JB said: "it was presented as a routine procedure". Neither remembered asking questions about the welfare of the animals that had been used in producing the antibodies. Nor did either remember asking about the cost/benefits of animals being used to produce material used for educational purposes.

Antibodies are not only a tool in the lab but are also one of the biggest growing drug markets today, especially for cancer treatment. You will encounter antibodies in academic research, company labs, and everyday life.

Questions

1. Do you remember when you first used antibodies in your studies? Was it in any way a memorable occasion?
2. When you first encountered antibodies, did you ask yourself where they came from or the circumstances under which they are produced?
3. Would you see this as an example of the WYSIATI heuristic in action? Why?
4. What does your reflection on this case tell you about ethical practices?

We will return to the looking at ethical questions around monoclonal antibody cultivation in a later chapter, when we consider the use of plastic in this process.

One of the features of heuristics that Kahneman points out is that, having jumped to a conclusion through the use of heuristics, our brain is generally very confident in the judgement or decision we have arrived at (as you saw above with the CRT, 83% of respondents did not reflect on their answer in at least one of the three questions). Hence, humans may not, under normal circumstances, be able to trust that we will spot predictable heuristic errors once we have made them. To address this issue, in ethical decision making, it is common to use particular *processes* or methods to force ourselves to slow down the decision making process. This is not only the case in ethical decision making – experts in any discipline typically work through a process in solving complex problems in their domain, although the process has often become so automated that, ironically, they choose the 'slowing down' process very quickly and sometimes without even being aware they are doing it. We will return to these processes in Chapter 5.

It is also important to be clear that heuristics are not necessarily bad: having mental shortcuts allows us to react quickly and without becoming paralyzed by the need to think through every decision. However, heuristics do sometimes lead to errors, and there are times when a slower, more reflective approach to decision making is appropriate.

Thus far, we have looked only at sources of bias which arise from the way an ethical question is framed or structured. Another factor that may influence our ethical judgement (or at least our ethical behaviour) is the people around us. We will consider this question in the next section.

How does social context affect ethical decision making?

On December 4th 2012, the New York Post newspaper carried a shocking and distressing image of a man in the seconds before his death. In the picture, Ki Suk Han, 58, was standing on a New York subway track, unable to climb back onto the platform, and watching the subway train which was racing towards him. Although the picture doesn't show them, there were people standing on the platform near Mr. Han, but none moved to help him. Seconds after the picture was taken, he was struck and killed by the subway train.

The photograph was taken by R. Umar Abbasi, a freelance photographer. Abbasi later said "The people who were standing close to him ... they could have moved and grabbed him and pulled him up. No one made an effort". Mr. Han had been drinking and was visibly drunk when an argument broke out between him and another subway passenger. But, when Mr Han was pushed and fell onto the tracks, no one acted to help him. (BBC, 2012; McKinley & Alani, 2017).

This is an example of what social psychologists refer to as the 'bystander effect', which describes the way in which a person's pro-social behaviour (that is, positive behaviour towards others) is influenced by the presence of other people. For example, one experiment looked at how students in a waiting room acted when they heard someone fall and call for help from an adjoining room. When the experimental subjects were alone in the waiting room around 75% responded. However, when they were in the room with another person who did not react to the noise (someone who was actually part of the experiment), only 7% went to help. When the experimenters put two experimental subjects who were strangers in together (rather than experimental subject with one researcher/actor who was told not to react), the pair helped in only 40% of cases (Latané and Rodin, 1969). The presence of bystanders inhibited

the willingness to help. Indeed, as the number of bystanders increased, the percentage of people who helped decreased, and people took longer to help (see Clarke, 2003; 58).

The 'bystander effect' is not straightforward and seems to depend on many different factors (such as whether a 'victim' is seen or only heard). Nonetheless, there does seem to be some evidence that the presence of others reduces helping behaviour, under at least some circumstances. One possible reason is a potential *diffusion of responsibility*. Put simply, while a sole bystander is 100% responsible for helping, two bystanders may each feel that their responsibility is less than 100%, while one person in three feels even less responsibility and so on. The presence of others may reduce responsibility to such a degree that something that is everyone's responsibility becomes no-ones responsibility.

Bystander effects can also be linked to perceptions of competence (Cramer et al., 1988). To give a personal example, while passing through the city of Madrid on a hot day in September 2024, I saw a man who I guessed was in his early 20s, who appeared unconscious and who was lying on the side of a footpath near the busy Atocha railway station. The vast majority of people were walking past him, and a person working in a nearby food truck who could see him on the ground was also ignoring him. I stopped walking to observe him, but did not approach him. In this situation I lacked a range of different competencies: I do not speak Spanish, I have no medical training that might allow me to assess if he needed help, and I did not know the neighbourhood. In addition, I didn't know the phone number for the Spanish emergency services and wouldn't have been able to speak to them even if I called them. I lacked multiple necessary competences. In cases like this, perceptions of competence can slow a pro-social response. (In my case, after a few moments of anxiously standing, I started to ask other passers-by if they spoke English or French, and, if they did, I asked if they

thought the young man was ok. A pair of women volunteered to go to the train station to ask for help. Eventually, to my shock, one passerby roughly grabbed the young man and dragged him to his feet while speaking to him in Spanish. Within a few minutes, the young man was buying water at the food truck, and was, apparently, well).

Case Vignette: What is your responsibility if you are concerned about scientific malpractice?

Etienne is a Master student working on a semester project in a lab. The semester project is contributing to a bigger project that is being managed by a post doc and also has inputs from two PhD students.

When Etienne presents some of his collected data, the post doc in charge of the project suggests that a number of data points be removed from the analysis. These are outliers, they suggest, which have probably come from an error in data collection.

Etienne is a little uneasy about removing these outliers as it remains unclear to him how they arose. But the amount of data being removed is quite small, and no one else in the team raises any concerns about removing them.

Questions

1. What do you think Etienne should do in such circumstances?
2. What role would Etienne's perception of competence (with respect to the PhD students and Post doc) play in that decision?
3. What guidance could Etienne get from EPFL ethical guidelines on appropriate actions in this case (Take a look at EPFL Rules on research integrity (Lex 3.3.2, especially, Article 7 https://www.epfl.ch/about/overview/wp-content/uploads/2019/09/3.3.2_principe_integrite_recherche_an.pdf)?
4. Are there particular ways a lab could be organised / particular lab practices that might increase the chance that Etienne would share concerns in this case, if that were justified?
5. If you were faced with this decision, having considered the context, what do you think would be an ethical course of action?

Perceptions of competence can slow a pro-social response, but may also speed up a pro-social response. Cramer et al., (1988) found that the presence of other bystanders did not reduce the probability that qualified nurses would act in a 'bystander' experiment, but did reduce the probability for untrained bystanders. In my example from Madrid, a pro-social response from

a passing nurse, or perhaps a police officer would probably not be slowed by the presence of bystanders.

It is not only the presence of people around us that may impact on ethical behaviour. In addition to the cognitive intuitions identified above, there is also evidence that people have *social intuitions* about other people, which can impact pro-social behaviour. There is substantial evidence that when people first meet (or even see a photo of) a person they don't already know, they very rapidly arrive at intuitive judgements about that person based on internalised social stereotypes. Psychologists identify that these rapid judgements are framed in two dimensions: perceptions of *warmth* (friendliness, trustworthiness, empathy, and kindness), and perceptions of *competence* (intelligence, power, efficacy, and skill). As Cuddy et al. (2011: 74) state:

> An important source of error in warmth and competence judgments stems from pervasive stereotypes based on others' race, gender, nationality, religion, profession, socioeconomic status, and similar social categories that influence whether we view another person (or another views us) as warm or cold, competent or incompetent. We may, therefore, make decisions about whom to trust, doubt, defend, attack, hire, or fire based on imperfect data.

Fisk (2018) draws on empirical data as to people's warmth and competence judgements both in the US and internationally. From this she identifies what stereotypes (generally based on age, and socio-economic factors) are associated with each quadrant of a 2X2 Warmth/Competence matrix across many countries. She also identifies persistent ethic/religious stereotypes found in data in the US. Her results are summarised in the table below.

Table 2:A. Warmth and Competence stereotypes in common across many countries, and ethnic and religious stereotypes in the US (adapted from Fisk, 2018: 2)

	Low Competence	***High Competence***

High Warmth	Common: Elderly, Disabled, Children US: Italians, Irish	Common: Citizens, Middle Class US: Americans, Canadians, Christians
Low Warmth	Common: Poor, Homeless, Immigrants US: Latinos, Africans, Muslims	Common: Rich, Professional, Technical Experts US: Asians, Jews, British, Germans

Fisk et al. (2002) identify that these stereotypes also affect sub groups within broader social groups. In their research in the US, for example, they found 'housewives' were rated as being low competence/high warmth, 'business women' were rated as high competence/high warmth and 'feminists' were rated as high competence/low warmth.

These rapid and intuitive social judgements can give rise to *implicit (*or *unconscious) bias* in interactions. Unconscious bias refers to stereotypes that a person themselves is not necessarily aware of (i.e. unconscious) that affect the expectations we have and the way we interact (i.e. bias) with people from particular social categories (Greenwald & Banaji, 1995). Simon et al. (2020), for example, studied the impact of social judgements on readiness to help a fellow student during a teamwork task. They found that both perceptions of warmth and competence predicted an increase in helping intentions. Hence, who we intuitively want to help may also be affected by social stereotypes.

Conclusion

In this chapter we have seen that while the human brain is capable of many amazing feats of judgement, one of the ways in which it achieves this is by using cognitive short cuts which allow it to quickly arrive at decisions. Often, these intuitive decisions are 'good enough', but

importantly, they can give rise to predictable, systematic errors of judgement. Some of these predictable errors of judgement include:

- An 'anchoring' effect, whereby a judgement is affected by the presence of some prior suggestion regarding the an outcome.
- A 'priming' effect, whereby a context which activates some prior memories influences a person's judgement or actions.
- A 'narrative over number' effect, whereby the presence of information that can be construed as a story influences a person's judgement.
- A 'what you see is all there is (WYSIATI)' heuristic, whereby the presence of some information leads us to assume we have all the information necessary to make a decision.

In addition, social dimensions can impact our intuitions and actions:

- the presence of others, as well as our judgements of our own competence, can impact on how likely people are to engage in pro-social behaviour
- rapid and unconscious intuitions arising from stereotyping, and these can also impact interaction and helping behaviour towards others.

Each of these sources of intuitions can produce biases in ethical decision making. Recognising these biases in ourselves can help us reduce their impact on our behaviour.

References

BBC (2012) New York Post photographer defiant over subway-death image. BBC 5 December 2012: https://www.bbc.com/news/world-us-canada-20611140

Clarke, D. (2003) *Prosocial and Antisocial behaviour* London: Routledge

Cramer, R. E., Mcmaster, M. R., Bartell, P. A., & Dragna, M. (1988). Subject competence and minimization of the bystander effect. *Journal of Applied Social Psychology*, *18*(13), 1133-1148.

Cuddy, A. J., Glick, P., & Beninger, A. (2011). The dynamics of warmth and competence judgments, and their outcomes in organizations. *Research in organizational behavior*, *31*, 73-98.

Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J. (2018). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. In *Social cognition* (pp. 162-214). Routledge.

Fiske, S. T. (2018). Stereotype content: Warmth and competence endure. *Current directions in psychological science*, *27*(2), 67-73.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, *19*(4), 25-42.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological review*, *102*(1), 4.

Kahneman, D. (2011). Fast and slow thinking. *Allen Lane and Penguin Books, New York*.

Latané, B. & Rodin, J. 81969) A lady in distress: inhibiting effects of friends and strangers on bystander intervention. *Journal of Experimental Social Psychology* 5, 189-202. https://doi.org/10.1016/0022-1031(69)90046-8

McKinley, J.C & Alani, H. (2017) Man Who Pushed Passenger to His Death on Subway Tracks Is Acquitted. *New York Times* 17 July 2017.

Simon, J. C., Styczynski, N., & Gutsell, J. N. (2020). Social perceptions of warmth and competence influence behavioral intentions and neural processing. *Cognitive, affective, & behavioral neuroscience*, *20*, 265-275.

Vohs, K. D., Mead, N. L., & Goode, M. R. (2006). The psychological consequences of money. *Science*, *314*(5802), 1154-1156.