UNIVERSITY OF SOUTHAMPTON

FACULTY OF PHYSICAL AND APPLIED SCIENCES

Electronics and Computer Science

Helping Users Adopt and Delegate Agency to Autonomous Agents in Everyday Life

by

Jhim Kiel M. Verame

Thesis for the degree of Doctor of Philosophy

August 2018

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
Electronics and Computer Science

Doctor of Philosophy

HELPING USERS ADOPT AND DELEGATE AGENCY TO AUTONOMOUS AGENTS IN EVERYDAY LIFE

by Jhim Kiel M. Verame

Autonomous agents are designed to take actions on behalf of users, acting autonomously upon data from sensors or online sources. However, the performance and actions of such agents are liable to uncertainties. As such, the design of interaction mechanisms that enable users to understand the operation of autonomous agents and flexibly delegate or regain control is an open challenge for HCI. Against this background, in this thesis we report on three studies designed to better understand how to help users interact with autonomous agents. In particular, we begin by understanding how people deal with uncertainties when delegating agency to autonomous services. We then examined the impact of different agent-designs and feedback mechanisms, inspired by the factors that encourages people to delegate agency in an everyday setting. Based on our findings, we discuss key implications for the design of future autonomous technologies and the design of interaction mechanisms to help users make the best use of such systems.

Contents

List of Figures

List of Tables

Declaration of Authorship

I, Jhim Kiel M. Verame, declare that the thesis entitled *Helping Users Adopt and Delegate Agency to Autonomous Agents in Everyday Life* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as: (Verame et al., 2016) and (Verame et al., 2018)

Signed:...

Date:..

Acknowledgements

First and foremost, I would like to express my deepest gratitude and appreciation to the excellent guidance provided by my supervisors, Enrico Costanza and Gopal Ramchurn. Without their constant push for success, like expecting me to submit a conference paper in my first month, I would not have been successful in my PhD.

I would also like to mention my friends and colleagues from the Agents, Interaction and Complexity Group during my time in Southampton. Thank you for your friendship and advice, especially during the first half of my PhD. My thanks also go to the members of the University College London Interaction Centre for welcoming me and for providing me a great new environment to continue my research when I moved to London.

I also acknowledge my funding source, the Engineering and Physical Sciences Research Council UK, that made this PhD possible. In particular, the studies I conducted were supported by the ORCHID (EP/I011587/1) and the AIoT (EP/N014243/1) projects.

I would also like to mention the love and moral support of my friends and family. I am most grateful for my parents, Jona and Abner; my brothers, Aj and Arjon; and my significant other, Daphne. Thank you for believing in me and for always being there for me, especially through the difficult times. Last but not the least, I thank the Father in Heaven for the many blessings He continues to provide to me.

Research Acknowledgements. Other than my supervisors, I was fortunate enough to work with a number of people throughout my PhD. I would like to take this opportunity to highlight their contributions in the research projects included in this thesis:

- Chapter 3 (veg box study): I collaborated with Dr. Joel Fischer, Prof. Andy Crabtree, Prof. Tom Rodden and Prof. Nick Jennings. They provided feedbacks throughout the various stages of the research project.
- Chapter 4 (groupbuying study): I worked with David Muehlebach. He was involved in the discussion of the research project at its early stages and helped conduct some of the experiments.

In addition, I thank my examiners, Prof. Matt Jones and Dr. Sebastian Stein, for their excellent suggestions on how to improve the thesis and for making the viva a memorable and delightful experience.

Chapter 1

Introduction

Computing systems are no longer scientific tools only used for industrial processes and in research labs. As was envisioned, we now live in an era where networked computers, smart phones, wearable gadgets and other ubiquitous computing devices are commonplace and have become prevalent in our daily activities (Weiser, 1991). Since these systems are often interconnected and equipped with different types of sensors, they are able to gather and generate vast amounts of data at unprecedented speeds. However, people have limited capacity to deal with such data deluge. As such, it has become necessary to transfer the responsibility, or *delegate agency*, over this data to *autonomous agents* (Jennings et al., 2014). In this thesis, we define autonomous agents as entities able to actively monitor and analyse data streams. However, unlike digital assistants, such as recommender or navigation systems, not only do they provide guidance or suggestions to help users make an informed decision, autonomous agents can make decisions and perform actions on our behalf. For example, a software that provides suggestions to users about what items to buy can only be considered as a recommender system. However, if it can also purchase on behalf of the user what to buy, then it is an autonomous agent. Autonomous agents are also proactive, in the sense that they are capable of initiating an action without the intervention of a human user. In addition, unlike automated software programs, autonomous agents have social capabilities, in that they are able to interact with other agents or users, and they use such interactions to improve its processes to achieve their goals. For instance, an OCR software may be considered an autonomous agent if it can utilise user input to improve its recognition accuracy. Practical real-world examples include appliances[1] that autonomously order part replacements based on usage to reduce delays and costs; smart thermostats[2] that take into account temperature and presence data, as well as explicit user input and the energy market to optimise the comfort of the occupants while minimising their costs; and *software agents* that can bid for users in on-line auction websites[3].

[1] e.g. https://infinity.brita.com/.

[2] http://www.nest.com/.

[3] http://www.snipeswipe.com/.

Generally, autonomous agents are based on techniques such as machine learning and artificial intelligence to process input data, be it from sensors or online sources, and automatically take decisions to guide their autonomous operations. While these agents are expected to work seamlessly, such expectations may not always be met in practice because of noise and biases in real-world data, limited size of training data sets, discrepancies between computationally feasible models and complex real-life systems. A real-world example of this is when the sensors of a self-driving car in the U.S. failed to distinguish another vehicle on the road, which then led to an accident [4]. Indeed, as depicted in this example, the results of automatic data analysis and classification may often be liable to considerable *uncertainty*. Such uncertainty may come from different sources, such as uncertainty in the environment of the agent; uncertainty in how the agent will cope with its environment; and uncertainty in how well the agent will perform its task. As such, decisions and actions of autonomous agents might be undesirable and may require human intervention. There may be situations where agents should be allowed to act completely autonomously, but at other times be guided by much closer human involvement (Scerri et al., 2002; Alan et al., 2016a). As a consequence, the design of interaction mechanisms that can help users deal with the uncertainty caused by delegating agency is an open challenge for HCI (Jennings et al., 2014; Verame et al., 2016, 2018). In this thesis, we are particularly interested in investigating agency delegation to *autonomous software agents*, given the lack of research around this topic.

1.1 Research Motivation

Research around interaction with autonomous agents for specialist applications (e.g. vehicle teleoperation or aviation) date back to the 1970s (Sheridan and Verplank, 1978). More recently, research has started to investigate autonomous agents in *non-specialist applications*, applications where users are not expected to be trained to use them. These include off-the-shelf autonomous products in the home, such as the Roomba (Forlizzi and DiSalvo, 2006; Sung et al., 2007) and the Nest thermostat (Yang and Newman, 2013; Yang et al., 2014). Other studies of non-specialist autonomous agents are based on deployed research prototypes for applications such as doing the laundry (Bourgeois et al., 2014; Costanza et al., 2014) and setting energy tariffs (Alan et al., 2016a,b). Findings from these studies have revealed that people are willing to delegate agency to autonomous agents. However, uncertainty about their capabilities and their quality of work has resulted in people not being able to appropriately delegate agency to these agents. Studies have shown that there are instances where users underutilise these agents. This is often because of the associated costs and risks involved when delegating control, such as worrying about losing money or causing an accident when the agent makes a mistake (Dzindolet, 2003; Alan et al., 2014). Other studies have shown that

[4] http://bit.ly/29byniC

people also tend to overly rely on autonomous agents. This bias towards autonomous capabilities occurs when users have little knowledge about how these agents work or when they assume that they work reliably at all times (Parasuraman and Manzey, 2010; Yang and Newman, 2013). Indeed, these studies suggest that people do not effectively delegate agency to autonomous agents and research has not fully exploited such issue. Addressing this issue is important, as it should also lead to a better adoption of autonomous agents (Yang and Newman, 2013; Yang et al., 2014).

Recent research has also outlined the challenges for designing interactions with smart products, such as the delegation of control between people and objects, and how to position objects in human activities as partners (Rozendaal, 2016; Cila et al., 2017). In particular, the partnership between users and agents has been considered essential for designing future autonomous agents (Klien et al., 2004; Farooq and Grudin, 2016), as poor partnership between users and agents can become increasingly costly and catastrophic (Lee and See, 2004). To enable this partnership, it is therefore important to explore the notion of agency delegation between users and agents. Identifying the factors that support and hinders agency delegation will enable us to formulate implications for the design of human-agent interactions. These interactions should also improve the *adoption* of autonomous agents, which refers to to the acceptance and inclination of people to use the technology. By so doing, people's utilisation of autonomous agents will improve, helping people take advantage of the capabilities of the agents.

Another motivation for this research is that autonomous agents are also yet to be explored in other mundane and everyday tasks, especially in situations where there are consequences to the users. In particular, no research so far has investigated user interaction with autonomous agents in domestic food practices. Food practices are particularly interesting as they offer a rich and challenging scenario for technological intervention because they are dependent on a multitude of contingent concerns, such as dietary requirements, different preferences of household members and food availability. Moreover, a number of our existing food practices already include some form of delegation to an autonomous entity, be it delegating cooking (e.g. take-aways) or food shopping (e.g. food box schemes). Based on this premise, a major part of this thesis aims to understand agency delegation to autonomous agents in the context of food-based practices.

1.2 Research Challenges

By undertaking this research, there are a number of research challenges we hope to address. The following outlines the key challenges we face in completing this research:

- **Designing acceptable agent behaviour** – As prior research suggests, an autonomous agent is more likely to be accepted and considered trustworthy if it displays personality characteristics similar to those of the user (Lee and See, 2004). A key research challenge then is understanding how to design autonomous agents such that they act in an acceptable manner to users to be considered trustworthy. Research has also shown that people respond socially to technology, and how people interact with computers are comparable to how they would interact with other people (Reeves and Nass, 1996; Milewski and Lewis, 1997; Nass et al., 1999). For instance, it has been shown that a system with autonomous capabilities that shows good etiquette (e.g. the system waiting for user to be free before sending an alert) improves a user's trust in that system (Parasuraman and Miller, 2004). However, it has also been suggested that having an overly human-like personality might make the agents more deceptive and misleading (Norman, 1994). Given this, it is important to explore what human-like attributes can be integrated into the design of an autonomous agent to improve its adoption.

- **Providing appropriate interaction mechanisms** – People have a natural tendency to either underutilise or at times rely too much on autonomous agents, which can cause undesirable results (Sutherland et al., 2015). These failures can cause user frustration and may eventually lead users to abandon the technology (Yang and Newman, 2013). It has been suggested that the root of this problem does not lie in the autonomous capabilities of such agents, but instead it is caused by inappropriate design (Norman, 1990). A design solution that has been advocated to address this issue is to make these systems more transparent, by providing meaningful explanations of their actions and outcomes (Tullio et al., 2007; Lim et al., 2009). However, providing explanations can also result in information overload, which can confuse and overwhelm users (Yang and Newman, 2013; Yang et al., 2015b). There is then a challenge in designing interactions that can encourage users to adopt non-specialist autonomous agents and also help users delegate agency to these agents, without increasing their mental workload.

- **Evaluating agency delegation** – It is important to study autonomous agents in a situation where delegating to such agents can potentially disrupt people in other activities, such as the case in the real world. These evaluations should then have *ecological validity*, which is defined as the extent that the observed behaviour in the evaluations can be generalised to a natural real-life behaviour (Soegaard and Dam, 2013). Because of the need for ecological validity to examine user adoption of autonomous agents, especially in non-specialist applications, evaluations in this area have typically been carried out as field trials. Although we agree with the importance of field trials, it is potentially difficult, or perhaps impossible, to compare interface features or agents' designs through this study method, because of the lack of control over external variables beyond our interests. A key challenge then is to develop alternative evaluation methods that account for ecological validity to study interactions with autonomous agents. We see the opportunity for lab studies as a potential alternative because they typically allow for precise measurements to compare alternative experimental conditions, and they also tend to be faster and cheaper to run than field trials.

In summary, there is a need to help users take advantage of the capabilities of autonomous agents, especially in situations where there are consequences to non-expert users. In our research, we hope to achieve this by understanding what characteristics should autonomous agents possess, and also by investigating what interaction mechanisms can help users make use of these agents. Based on the motivations and challenges of this research, we explicitly define the research questions of this thesis in the next section.

1.3 Research Questions

The summary of the main research question of this thesis is as follows: *how should an autonomous agent be designed for people to willingly adopt and delegate agency to it?*. In particular, we focus on understanding interactions with autonomous agents in non-specialist applications. We tackle this research question by addressing the following sub-questions:

- R1: How do people deal with the uncertainties of existing services that exemplify agency delegation, particularly in everyday situations?
- R2: What *characteristics* should autonomous agents have to help people adopt them despite their uncertainties?
- R3: What *feedback mechanism* can help people to effectively delegate agency to autonomous agents that are subject to uncertainties, such that people do not rely too much or too little on them?

  Finally, an additional research question we wanted to address is related to the experimental design of the user studies:

- R4: How can we design user studies to compare alternative agent-designs in the lab with *ecological validity*?

In the next section, we detail how this thesis has been structured in order to address these research questions.

1.4 Research Structure

To be able to tackle the challenges stated, this thesis has been structured by firstly investigating a wide perspective on people's agency delegation in everyday practices. In particular, to address R1, in the first study outlined in Chapter 3, we focused on understanding why is it that people are already willing to hand over control to services that act *autonomously*, even though such services are bound to a certain degree of uncertainty. Applying the knowledge obtained in this exploration, we then studied a number of concrete applications, by looking at user interactions with autonomous software agents. In particular, the implications of the first study were implemented and tested in Chapter 4. This was done to address R2, which is to investigate whether the characteristics that enabled people to delegate agency to an autonomous service can also help people delegate agency to an actual software agent. Undertaking the second study then resulted in an implication of what type of feedback mechanism can be used to engage people to interact with autonomous agents. This feedback mechanism was then evaluated as detailed in Chapter 5, to address R3.

The structure of this thesis is further detailed in the following section.

1.5 Research Overview

This thesis is divided into a number of chapters as follows:

- Chapter 2 provides the necessary literature review that has influenced this research. The chapter begins by describing existing studies around user interaction with autonomous agents, both in specialist and non-specialist applications. We also reviewed proposed interaction techniques to help users make sense of autonomous agents. Finally, we briefly reviewed research from the field of Behavioural Economics to inspire the study methods used in some of our evaluations, and to also provide us a perspective about how people make decisions under uncertainty about the world and incentives.

- Chapter 3 explores the notion of agency delegation in the everyday practice of food. In particular, this chapter presents a mixed qualitative study looking at an established service that exemplifies agency delegation as a way of understanding how people deal with uncertainty. As an instance of agency delegation in everyday life and related to food practices, we study a *vegetable box (veg box) scheme*[5] and the food practices around it. In subscribing to the veg box scheme users delegate agency to the service provider as to the particular food items they receive in much the same way as one might delegate agency to an autonomous service, such as a thermostat to which one might delegate the decision to heat or cool a home. While participants in one group

[5] e.g. http://www.riverford.co.uk/

were existing subscribers of a veg box scheme, the participants in another group were subscribed to the scheme for the duration of the study. Through this arrangement, we could contrast existing and emerging practices around agency delegation, including the disruption of existing food routines. A number of design implications are discussed, such as supporting user creativity and taking into account personal values to warrant agency delegation.

- Chapter 4 further examines how designing autonomous agents that take into account personal values can affect people's inclination to adopt the technology. In particular, we looked at the delegation of agency to shopping agents, especially in group buying situations, where people find others to share the costs and the contents of a wholesale product (e.g. a pack of rice containing 6 individual units). In this chapter, we report on a controlled lab study designed to test whether the notion of fairness, a personal human value, can impact one's inclination to make use of an autonomous group buying agent. The chapter provides design implications for future autonomous agents, including a design solution for dealing with uncertainty through ensuring that an agent's performance is clearly conveyed to users.

- Chapter 5 further explores another design solution to help users deal with uncertainty. In more detail, this chapter reports on a lab study designed to investigate whether displaying the confidence information of an autonomous agent can improve its adoption and help user appropriately delegate to it.

- Chapter 6 summarises and discusses the achievements of this thesis. In this chapter, we start by summarising each study and discussing their implications. We then describe future research opportunities for investigating interaction with autonomous agents based on the work presented in this thesis.

  It is worth noting that the studies in this thesis have been presented in a different order to how they were originally conducted. This was done to express a clearer approach to the problem stated in this thesis. In the interest of transparency, the order in which the studies have been conducted is as follows: the work presented in Chapter 5 was completed first, followed by the veg box study in Chapter 3, and the groupbuying study described in Chapter 4 was the last study completed.

1.6 Research Contributions

The following outlines the key contributions of this thesis, based on the findings of our studies.

- We present our findings from a qualitative study describing how people deal with uncertainty around an established food-related service that exemplifies agency delegation (Chapter 3), which addresses R1. Our findings suggest that agency delegation must be warranted, that it must be possible to incorporate delegated decisions into everyday activities, and that delegation is subject to constraint.

- We proposed a number of recommendations for designing future autonomous agents, especially for agents supporting food-based practices to address R2. For instance, we outlined the importance of designing agents that support human values, such as enabling creativity (Chapter 3) and fairness (Chapter 4), to encourage users in adopting autonomous agents despite their uncertainties.

- We provide design solutions to help users interact with autonomous agents, addressing R3. First, we suggest that if agents are designed to support people's values (e.g. helping people save energy), information relevant to these values should be easily available to the users (e.g. provide data about financial savings on utility bills) (Chapter 3). Having this information will help users decide whether to keep delegating agency or to revoke it. Second, our findings suggest that rather than displaying an explanation of how an autonomous agent works, designers should instead focus on making the capability of the agent clearer, so as to help users set an appropriate expectation on its outcome (Chapter 4). Finally, we showed that displaying the certainty of an autonomous agent about the quality of its work, which we call its *confidence information*, can also increase the adoption of the agent and help people appropriately delegate agency to it (Chapter 5).

- To resolve R4, we developed a study method that can be used for evaluating the effectiveness of interaction mechanisms designed to support agency delegation (Chapter 4 and 5). Specifically, using our translation of behavioural economics research methods, we incorporate experimental financial incentives and carefully designed tasks to maintain high degree of ecological validity in running controlled lab studies.

In addition to the above contributions, the following paper describes the work presented in Chapter 3:

Verame, J.K.M., Costanza E., Fischer J., Crabtree A., Ramchurn S.D., Jennings N. and Rodden T. "Learning from the Veg Box: Designing Unpredictability in Agency Delegation". In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18). Montreal, Canada. Paper No. 447.*

A version of the content in Chapter 4 will be submitted to CHI:

Verame, J.K.M., Muhlebach D., Costanza E., Ramchurn S.D. and Jennings N., "Designing Autonomous Agents for Online Grocery Shopping". (to be submitted)

The work presented in Chapter 5 was published in the following paper:

Verame, J.K.M., Costanza E., and Ramchurn S.D. "The Effect of Displaying System Confidence Information on the Usage of Autonomous Systems for Non-specialist Applications: A Lab Study". In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16). San Jose, USA, 4908–4920.*

In parallel to the work reported in this thesis, throughout my PhD, I contributed to other research related to interaction with smart and autonomous systems, which resulted in the following peer-reviewed publications for which I am a co-author:

Garcia Garcia, P., Costanza E., Ramchurn S.D., and Verame, J.K.M. "The potential of physical motion cues: changing people's perception of robots' performance". In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16). Heidelberg, Germany, 510–518.*

Garcia Garcia, P., Costanza E., Nowacka D., Verame, J.K.M. and Ramchurn S.D."Seeing (Movement) is Believing: the Effect of Motion on Perception of Automatic Systems Performance". In *Human-Computer Interaction (2018): 1-51.*

Additionally, one other related piece of work will be submitted to CHI:

Nowacka D., Yang R., Costanza E., Verame, J.K.M., Fischer J., Ramchurn S.D. and Jennings N., "A Field Trial of Autonomous Food Shopping". (to be submitted)

Chapter 2

Literature Review

This chapter presents the essential literature review that has influenced this research. In particular, we identify the issues and challenges of this research through understanding the use of autonomous agents in specialist applications, how people respond to agent-based decision support systems and people's orientation and experience of autonomous agents in everyday situations. We also cover existing work around domestic food practices, especially in terms of potential opportunities for technological interventions and understanding agency delegation in this application area. Finally, we provide an overview of behavioural economics literature, firstly to inform the design of study methods that can be used to evaluate user interaction with autonomous agents in lab setting; and secondly, to understand people's decision-making in situations where their decisions have financial implications.

2.1 Interactions with Autonomous Agents

In this section, we provide an overview of studies that involve users interacting with autonomous agents. In particular, the review of these studies is divided into three main categories: specialist autonomous agents, agent-based decision support systems and autonomous agents in specialist applications. We conclude the section by also reviewing current design approaches to help users interact with autonomous agents.

2.1.1 Specialist Autonomous Agents

Numerous studies of users interacting with autonomous agents have been conducted in specialist applications, such as in the military (de Tjerk et al., 2010; Helldin et al., 2014) or aviation (Parasuraman et al., 1993; Dehais et al., 2015), for which users would need and receive considerable amount of training. These studies typically examined the effect of increased autonomy on users' cognitive workload and how that impacts their performance in completing specific tasks through lab studies. For example, researchers have focused on the operation of agent teams by single or multiple users in search tasks (e.g. finding victims in a simulated disaster event) (Goodrich et al., 2007; Lewis et al., 2010; Wang et al., 2006). In more detail, participants either operated a team of manually operated agents or monitored a team of autonomously moving agents. Other work has investigated how different autonomy levels affect cognitive workload and how that impacts user performance of allocating tasks to multiple (virtual) unmanned aerial vehicles (UAVs), all of which are represented by software agents (Lin and Goodrich, 2015; Ramchurn et al., 2015). Participants performed the search tasks either through manual or mixed-initiative task allocation of UAVs. More specifically, in the manual condition, participants had to manually assign tasks for all agents. In the mixed-initiative condition, participants received recommendations on how to allocate tasks for the agents, which could either be accepted or rejected. Results from these studies showed that higher autonomy improved user accuracy with the tasks and reduced cognitive workload.

These studies assumed that users would interact with autonomous agents, enforcing agency delegation. In contrast, our work focuses on non-specialist activities where users would not normally be trained to use the agent and we examine whether users would adopt the autonomous agent or not. As such, we also look at situations where the use of a system with autonomy can be ignored but with the consequence of increasing workload. This is the case with agent-based decision support systems, which we detail in the next section.

2.1.2 Responding to Agent-based Decision Support Systems

Decision support systems, which can also be considered as instructional agents or recommender systems, are agent-based systems that provide users with suggestions or recommendations to help them complete a task. Users have the freedom to choose whether to follow or ignore the suggestion provided by an agent. For instance, a number of studies have looked at how people react to suggestions from an agent through games that uses real world locations in game play, so-called *pervasive games* (Moran et al., 2013; Jiang et al., 2014). These studies focused on how team members collaborate amongst each other to act upon suggestions received from an agent. Findings from these studies suggest that while human teams are typically inclined to follow suggestions from an agent, suggestions from an agent are likely to be rejected if the users are engaged with other tasks or if previous suggestions by an agent failed to achieve users' expected outcome.

Other studies have looked at users' reliance on decision support systems. In the study by Smith et al. (1997), they focused on how expert users (pilots and dispatchers) makes use of a software agent for aircraft flight planning. The task of the participants was to plan the best route for an aircraft in terms of fuel consumption and time efficiency.

Each participant was placed in one of the three system design: *sketch-only* (S), *sketch and route constraint* (SR), *automatic route-constraint and sketch* (A). In S, participants specified the flight path and the agent would calculate the estimated time and fuel consumption of such path. In SR, participants could place constraints and the agent would create a flight path on behalf of the user based on such constraints. Lastly, in A, the agent would automatically create a suggested a flight path based on default constraints. In both SR and A, participants were able to review the flight path created by the agent. The results of their study suggests that participants were overly reliant on the suggestions by the agent, especially when suggestions were presented before they attempt to create their own desired flight path. In some cases, participants would even choose to not evaluate what the agent suggested and simply accept them. In short, participants favoured the capability of the automation over their own, a concept known as *automation bias* (Parasuraman and Manzey, 2010).

Dzindolet (2003) examined how people's trust affect their reliance on decision support systems. The task for the participants was to go through a set of images and identify whether or not there is a camouflaged soldier in each image under a set time. Once a participant provided an answer for an image, the system would provide a suggestion that may or may not be different from the participant's answer. The participant is then given an opportunity to change the initial answer. All participants were told that the system is not perfect and may make incorrect suggestions. Results of their study showed that participants took the suggestions of the aid when there were no associated risks or costs. However, when participants were given monetary rewards for correct answers, they became less inclined to take the system's suggestions. Interestingly, when participants were provided an explanation of why errors may happen, they became inclined to rely on the system regardless of its accuracy and the monetary rewards. This can be problematic because if the accuracy of the system is poor, people might trust in it, which can lead to undesired results. In a similar vein, Sutherland et al. (2015) conducted a study to investigate how the cost of using automation affects users' reliance on decision support systems. In the study, participants played a game where they were required to create an army of orcs. To do so, they had to visit various locations in the game and choose the best orc in each of those locations. The orcs in each location varied in size, but the size differences were not always clear so at times, it would be difficult to determine the best orc. To help participants make a decision, a decision-support system that recommended the best orc was available. Using the system or failing to choose the best orc had a cost in the form of time (e.g. unable to move in the game for 30 seconds) to ensure real consequences in decisions made. The cost was varied during the study. Results showed that participants tended to overly rely on the automated aid when the cost of using the system was very low, even though there were clear differences between the sizes of the orcs. Participants also underutilised the system when the cost of using it was very high, even if using it would improve their decision and even in situations where using it was required.

These studies highlight that users tend not to correctly rely on suggestions provided by autonomous agents. These studies highlight that people have difficulties

When the cost of using an agent is relatively low, users tend to over rely on it. In contrast, users are less likely to rely on an agent when the cost of using it is high (Dzindolet, 2003). In our work, we are interested in designing interaction mechanisms that can address this issue and help people utilise autonomous agents, such that they "rationally" rely on them by having a clear understanding of the expected utility of the agents. Moreover, we focus on delegating agency to autonomous agents in a situation that involves multiple activities, much like in real life.

2.1.3 Autonomous Agents in Everyday Situations

Researchers in the HCI and Ubicomp communities have explored the adoption of emerging autonomous agents in the home environment. For example, Rodden et al. (2013) used animated sketches to solicit views from people about current and future agent-based energy systems. Participants were strongly against the idea of losing autonomy and control over their appliances in the home and generally felt that they could not trust autonomous agents to do their work for them. Similarly, Ball and Callaghan (2011) conducted an online survey showing the concept of *adjustable autonomy* in intelligent systems of a smart home. In this context, adjustable (or flexible) autonomy refers to the ability that users can dynamically change the level of autonomy of an agent at any given time. More specifically, this study focused on understanding what level of autonomy would people prefer for future smart systems. Responses show the diversity in how people accept autonomy and the majority have expressed that they would not prefer a fixed autonomy for all sub-components of an intelligent system. An example is that for personal use such as entertainment, people would prefer lower levels of autonomy, but higher levels of autonomy would be acceptable for controlling heating and lighting. Some people also felt that perhaps higher levels of autonomy would be good to begin with for an intelligent system, then lowering this autonomy over time as they become more accustomed.

The above studies elicited views from participants through mock-ups and scenarios. Other studies instead evaluated potential agent-based systems through actual usage of prototypes, such as for laundry management (Bourgeois et al., 2014; Costanza et al., 2014), smart energy systems (Alan et al., 2016a,b) and self-driving cars (Lee et al., 2014; Rödel et al., 2014). For instance, Alan et al. (2016a) conducted a number of field trials to study how a tariff-switching agent will be used by people in a situation where there are dynamically changing priced energy tariffs. In particular, they study the notion of *flexible autonomy*. Participants could choose from three settings: 1) the agent could suggest a tariff for them, but they would manually set the tariff; 2) the agent could set the tariff on their behalf and notify them once it is set and 3) the agent could fully autonomously set the tariff on their behalf. Their findings suggest that users were ready to delegate control to the agent to switch their tariff but were keen to monitor the performance of the agent. None of their participants felt comfortable leaving the agent to be fully autonomous, aligning with the results of prior work (Ball and Callaghan, 2011; Rodden et al., 2013). In a similar vein, Costanza et al. (2014) conducted a field trial to understand how people use energy in the home in a future scenario where electricity prices change in real-time. In particular, they developed an agent-based interactive system that is aimed at helping people do their laundry. Participants in this simulated scenario were asked to use the system to book a timeslot of when they would like to do their laundry. In the background, then, an autonomous agent would attempt to minimise the cost of the wash by charging a battery when electricity is cheap. Their findings highlight that some participants were able to integrate such a system into their existing laundry practices, but uncertainties and contingencies in everyday life hindered others from adopting the system.

The above findings suggesting that users tend to be inclined to adopt autonomous agents and integrate them in their day-to-day routines also emerge from studies of people using off-the-shelf autonomous products, such as the Nest thermostat (Yang and Newman, 2013; Yang et al., 2014) and the Roomba self-cleaning robots (Forlizzi and DiSalvo, 2006; Sung et al., 2007). Moreover, people develop strategies to work with these systems to make them deliver desirable performance. For example, users of self-cleaning robot develop routines, such as moving chairs around and raising the wires off the floor so that the robots could navigate through their house without getting stuck or choking up wires (Forlizzi and DiSalvo, 2006). Similarly, users of self-driving cars also develop strategies to alleviate the imperfections of such autonomous agents (Brown and Laurier, 2017). In particular, they deal with uncertainties by becoming aware and anticipating situations that can lead to errors, such as the presence of sunshine or approaching blind corners. However, research suggests that when users find it difficult to recognise how well an autonomous agent works, they become unable to change its behaviour when it does not work as expected (Yang and Newman, 2013). As a result, users become frustrated and their interaction with such systems decreases over time, which may potentially lead to distrust and even the abandonment of the technology (Lim et al., 2009; Lazar et al., 2015).

In summary, while studies have shown that people are willing to use autonomous agents, there appears to be an imbalance between user and agent autonomy, caused by the uncertainty of the autonomous agents. Users often create strategies to deal with the uncertainty. However, when they lack knowledge of how an autonomous agent works, they either disengage with the system (Yang and Newman, 2013), or they continuously monitor its actions (Alan et al., 2014), which adds workload to users. It is therefore important to establish ways to help users know when to grant or revoke agency from these systems, without the need for them to understand the technical details of how autonomous agents work.

2.1.4 Improving Interactions with Autonomous Agents

Researchers have explored different ways of improving how users interact with autonomous agents. For instance, based on research suggesting that people inherently follow human social norms such as empathy or politeness towards technology (Nass et al., 1999; Reeves and Nass, 1996), some research has suggested that agents should be designed with human-like attributes to encourage users to interact with them (Friedman and Kahn, 1992; Friedman et al., 2008; Liu et al., 2008). For example, a large body of work has focused on design approaches to make autonomous agents, particularly robots, human-like (Fink, 2012). By so doing, these agents become more familiar, explainable and predictable to users, encouraging users to interact with them. However, there are also some studies suggesting people do not always prefer their agents to act like them (de Melo et al., 2016; Basu et al., 2017). For instance, in the study by Basu et al. (2017), 15 participants took part in a driving simulation experiment to understand people's preference when it comes to the driving style of autonomous cars. Before the study began, participants completed a training phase to familiarise themselves with the simulator. They also experienced different driving styles and had to distinguish which one matches their own personal driving style. In the study, the car performed a number of specific driving tasks during the simulation, such as changing lanes or turning right on a green light. The car performed the manoeuvres based on a variety of driving styles, which is based on different levels of defensiveness (e.g. distance to other cars the larger the distance, the more defensive the driving). The study followed a think-aloud protocol, where participants were asked to express their emotions and feelings loudly as they experience the autonomous driving. Findings from their study suggest that participants preferred autonomous cars to drive at lower speeds than their normal driving style. One of the main reasons participants mentioned was to ensure safety, especially since they do not have (manual) control over the car. However, despite such studies, no research so far has explored how designing an agent with human-like attributes affects its adoption and whether people delegate agency to it, especially in situations where the agent can be ignored.

While applying human-like attributes to an autonomous agent may help people to interact with it, it is not always feasible to apply human-like attributes to agent design. Indeed, it is not the only solution to design effective human-agent interaction (Fink, 2012). To improve users' uptake of autonomous agents and to address expectation mismatch, other researchers have suggested that the intelligibility of such systems be increased. Prior work has suggested that this can be achieved through increasing system transparency, which refers to the explanation of actions and processes performed by a system (Helldin et al., 2014; Rodden et al., 2013; Yang et al., 2014, 2015b). For example, Lim et al. (2009) ran two experiments to analyse the effect of meaningful explanations describing *why* and *why not* a context-aware application behaved in a certain way. Their findings suggest that having these explanations improves users' understanding of the behaviour of the system and also increases users' trust in the system. However, studies also suggest that increasing system transparency may not always be a good solution (Helldin et al., 2014; Yang et al., 2015b). Novice users or those that are less technically minded may find such information difficult to understand. Moreover, reviewing this information requires plenty of time from the users, which increases their workload.

Other work has examined how displaying *confidence information* can be employed to resolve this problem (Beller et al., 2013; Helldin et al., 2013; McGuirl and Sarter, 2006; Desai et al., 2013). Confidence information refers to the estimated probability that an inference made by a system is correct, provided that the system has the correct model to interpret its data. Displaying confidence information has mostly been researched with an aim to find out how users interact with context-aware systems. For example, Lemelson et al. (2008) compared different visualisation methods of confidence information for location-based services, while Antifakos et al. (2005) asked participants to report whether they would check the settings automatically set up by a context-aware system given that they were in certain scenarios with different criticalities (e.g. while eating at a restaurant or while driving). Their findings suggest that participants were more willing to review the settings given that they were shown a display of confidence information, especially when the system's confidence level was low. Findings from a study by Lim and Dey (2011) also suggest that displaying the confidence information of context-aware systems can affect users' understanding and impression of such systems in a variety of ways. A user's understanding and impression of a system can be improved when it has mostly high confidence levels. However, displaying confidence information can be harmful in situations where the system has mostly low confidence levels, as users tend to lose trust in its capabilities. In contrast to these studies, our work focuses on the use of a functioning system to observe how users interact with autonomous agents rather than results elicited through reports of subjective preferences.

There have been studies examining confidence information using functioning prototypes. For instance, Antifakos et al. (2004) examined whether displaying confidence information can improve the usage of context-aware memory aids. Their results suggest that users do perform better with the display of confidence, especially when the confidence level is high. Similarly, the findings of Dearman et al. (2007) suggest that displaying confidence information can improve user performance in a search task using a location-based service application. In contrast to both studies, Rukzio et al. (2006) found that displaying the confidence information of an automatic form filler slowed down users and caused them to make more errors as they often double-checked fields with lower confidence levels.

Instead of focusing solely on performance, our work focuses on agency delegation and acceptance of autonomous agents.

Prior work has also investigated the effect of displaying the confidence information of self-driving cars (Beller et al., 2013; Helldin et al., 2013). Results from these studies showed that displaying the confidence information reduced the time it took for drivers to take control of a self-driving car and allowed drivers to spend more time not looking at the road. In a study by McGuirl and Sarter (2006), pilots were asked to complete a series of simulated flight exercises, requiring them to operate a number of manual tasks and monitor an automated system that would require users to take control at times. Results of this study indicate that pilots shown constantly updating confidence information were able to complete the flight tasks without failures and were better at estimating the accuracy of the automated system than pilots with only information about the overall reliability of the automated system (i.e. the accuracy of the automated system). In all three studies, the participants' choice to use the autonomous agent or not had no tangible consequence on them, which lowers the ecological validity of the studies. In our work, we aim to study interaction with autonomous agents through study methods with ecological validity.

One study has examined the effect of confidence information in an ecologically valid set up (Desai et al., 2013). In more detail, this study investigated the effects of displaying the confidence of a moving robot that could move fully autonomously or in semi-autonomous mode (i.e. users can control the direction of its movement). Participants had to monitor a moving robot, help it pass obstacles and occasionally complete a secondary task (clicking a circle on the screen). In their study, participants were rewarded based on task performance. The results showed that participants with the confidence information switched between full and semi-autonomy mode more than participants without the information. Particularly, participants were found switching to semi-autonomous mode whenever there is a drop in confidence, even though reliability did not change. Contrary to this work, we focus on users interacting with autonomous agents, rather than physically embodied autonomous robots. Furthermore, this study required participants to actively monitor the autonomous robot, which enforced the interaction. In contrast, our research is based on situations where people can choose to completely ignore an autonomous agent, because in reality, people can choose to not use it at all.

2.2 Food-based Practices

Studies of users interacting with autonomous agents highlight the opportunities and the importance of considering how such technologies can affect our everyday lives. However, none of those studies have looked at user interaction with autonomous agents in the food domain. The domain of food is a particularly interesting application for autonomous agents, as a number of our existing food practices already include some form of agency delegation, be it delegating cooking (e.g. take-aways) or food choices (e.g. food box schemes). As such, the majority of our work, particularly in Chapters 3 and 4, focuses on understanding how autonomous technologies can be developed and designed to support food-based practices. As such, we review prior research that has studied food-based practices, especially studies that focus on the potential of technological interventions in these practices. We also provide background around veg box schemes and group buying, which are central to our work presented in Chapters 3 and 4.

2.2.1 Food HCI

Researchers in HCI and Ubicomp have developed and evaluated technologies to support domestic food practices. For example, through lab studies, augmented kitchens have been evaluated to test their effectiveness to promote healthy cooking (Chen et al., 2010; Chi et al., 2008). Similarly, Reitberger et al. (2014) developed and evaluated Nutriflect, a system which informs users about their shopping behaviour through a situated display in the kitchen based on past shopping data. Findings from these studies suggest that displaying user food habits can raise awareness towards healthier food shopping and consumption. Instead, other research has focused on supporting cooking practices outside of healthy eating, such as helping users to execute complicated cooking tasks (e.g. cooking several recipes in parallel) (Hamada et al., 2005) and recommending easy-to-prepare meals to users (Palay and Newman, 2009). Other researchers have also looked at how technology can be used to support mealtime interaction and experience (Ferdous et al., 2016). While these works focused on the development of technical prototypes to support existing food practices, we focus instead on delegating agency to autonomous agents that support food-based practices.

Previous work on food-related research in HCI has examined existing food practices in the home environment. Cha et al. (2015) conducted an observational study which focused on understanding users' organisational habits in the kitchen (e.g. unpacking and storing of groceries) to formulate design implications for organisational robots. Similarly, Comber et al. (2013) conducted contextual inquiries with ten households to identify people's general domestic food practices, focusing on food purchasing and consumption. In contrast, Kuznetsov et al. (2016) conducted an in-situ fieldwork to observe the work of practitioners of at-home food science, which involves practices such as food preservation and fermentation. Their contribution focuses on how technologies can help practitioners to adopt food science as a habitual and everyday practice. Instead, Grimes and Harper investigated how user experience of eating and preparing food can be improved by designing technologies that *celebrate* positive user interactions with and around food (Grimes and Harper, 2008). Hupfeld and Rodden (2012) conducted a study focusing on users' dining practices, particularly in the interactions around the dining table. Based on their findings, they formulate implications for technologies seeking to augment eating experiences. In a similar vein, Clear et al. (2013) conducted a study using various enquiry techniques to capture students' food preparation activities, focusing on their cooking habits and greenhouse gas emissions. They discuss potential design interventions for more sustainable cooking practices. Prior work has also investigated how generational homemade family cookbooks are used in existing cooking traditions and practices through various ethnographic techniques (Davis et al., 2014). Their contribution focuses on the design implications of digital homemade cookbooks to support the retention and heritage of family recipes. In contrast, other studies examine users' experiences of food waste, and its connection to other food practices and reasoning behind them (Ganglbauer et al., 2013, 2015).

Our work is similar to the studies described above as we are also interested in food-based practices. However, our work extends current research as we focus in delegating the agency of food-based practices. In particular, we achieve this in two ways: firstly, in Chapter 3, we focus on how users embed a food service that exemplifies agency delegation — a veg box scheme — in their existing practices around food, including shopping and food preparation; and secondly, in Chapter 4, we investigate interaction of users with autonomous software agents that are designed to help them in grocery shopping, particularly in group buying situations. In the following subsections, we give a brief background about veg box schemes, detail existing studies around such services, and also provide an overview of research around group buying.

2.2.2 Veg Box Schemes

Veg box schemes historically started as a way for consumers to support local farms and get fresh fruit and vegetables directly from them (Petrini, 2007); originally consumers would pay upfront for a yearly or half-yearly supply of produce, which they would then get delivered (or collect) weekly or fortnightly. Nowadays most schemes have taken a more commercial turn, and those who run the scheme may grow only part of the produce that they distribute to customer, with the rest being sourced from other farms that may be local or not. In particular, in the UK, where our second study took place, fruits in some periods of the year tend to be imported from abroad. Subscriptions are no longer over the long term, but customers can buy boxes on an individual basis, generally just with the constraint of a minimum size order. Some schemes allow customers to select exactly what produce they want to receive or to specify blacklist of items that they never want to receive. Most services advertise the list of items in the weekly box a few days in advance, either on websites or via email, and many also publish recipes as suggestions for how to use the box content.

More recently some veg box providers, as well as other companies[1], also started to offer *recipe box subscriptions* – a weekly or fortnightly box that contains the exact ingredients for cooking a specific number of predefined recipes. In some cases, the main ingredients are selected from what is locally in season. Even though there are some similarities between recipe boxes and veg boxes (e.g. the periodic delivery, and the limited control over what gets delivered), a critical difference is that recipe boxes are self-contained, and do not require, nor allow, integration with customers' existing food practices in the same way that veg boxes do (as the content of the veg box generally needs to be complemented by other items in order to be turned into meals). For this reason, we decided to focus strictly on veg boxes, as we believe they offer a richer opportunity to observe whether and how such an autonomous service rubs against existing everyday practices.

Studies around sustainable food consumption practices often looked at veg box schemes. Typically, such studies conduct surveys which involves asking veg box users to identify motivations and experiences consuming local organic food (Seyfang, 2006). For example, Brown et al. (2009) conducted surveys in France and England to determine the motivations and barriers for consumers to subscribe to fruit and veg box schemes. Participants were asked to rank a number of different possible motivations and barriers based on their preferences. Findings from their study suggest that the box's low food travel mileage is the most important motives for English veg box consumers to subscribe (*altruistic reasoning*). In contrast, French consumers' most important motive is receiving high quality produce (*hedonistic reasoning*). Consuming out-of-season food was reported as the main barrier for English consumers to commit to a veg box scheme. The next main barrier for English consumers was identified to be the cost of the box, which is also the main barrier for French consumers. There are also other studies of veg box schemes outside motivations to subscribe, such as one that focuses on evaluating the logistics performance of existing veg box scheme services (Bosona et al., 2011). In contrast to these studies, our work, especially in Chapter 3, focuses on people's actual usage and adoption of the veg box, going beyond understanding consumer motivations. Most importantly, to our knowledge, veg box schemes have never been studied as an instance of a service exemplifying agency delegation – we are particularly interested in how consumers deal with the autonomous decision-making element in veg box schemes, i.e., the lack of user control over the items included in the box, and also how consumers go about the uncertainties associated with veg box schemes.

The veg box scheme portrays a particular situation of uncertainty, one where users are uncertain about the agent's decision on the contents of the veg box. Here, uncertainty is associated with an individual agent, particularly in its lack of information about the precise number of available produce. In other situations, uncertainty might be a result of an agent interacting with other agents, which is the case in multi-agent systems.

[1] e.g. https://www.hellofresh.com

For instance, in multi-agent coalition formation, there is uncertainty associated with whether an agent can successfully form a group with other agents, in order to achieve a goal. A real-world example of this would be in group buying situations, where buyers come together to share the cost and contents of a wholesale product. Based on this, we are interested in exploring how people deal with uncertainty as a result of group interactions. We further provide details of group buying in the next section.

2.2.3 Group Buying

The core concept of group buying is to gather enough people to achieve a volume of products to create the basis for a reduced purchasing price (Kauffman et al., 2010). In traditional group buying, people form a group of potential customers to buy goods or services, ranging from household items to cars. These groups are mostly formed online and with combined buying power, unprecedented bargains are possible. These collective group buying auctions have the chance to outperform recommended retail prices when the number of consumers in a group is large enough (P.K. Kannan, 2001). However, group buying also comes with its disadvantages. It can be time-consuming to find a group to join and the process can be quite complex for a user to understand, which can lead to people not taking advantage of it (Kauffman and Wang, 2002; Lim, 2003). We see the potential, then, for autonomous agents to help in this process, thereby increasing potential consumer participation in group buying. Indeed, while prior work has looked at how autonomous group buying agents can be developed (Lee et al., 2002; Asselin and Chaib-draa, 2002; Ito et al., 2002), we focus instead on studying whether users would adopt autonomous group buying agents and how users would interact with them.

2.3 Behavioural Economics

Our work involves the delegation of agency to autonomous agents in situations where such delegation has financial implications. Based on this premise, we turn to behavioural economics for inspiration to inform our research. Behavioural economics is concerned with the effects of psychological factors on people's economic decisions (Kahneman, 2011). We are particularly interested in applying behavioural economics research methods to understand agency delegation in the lab without compromising ecological validity. Furthermore, as we are interested in employing autonomous agents in bargaining scenarios, we reviewed existing studies of how people act in such scenarios. We provide details of these in the following subsections.

2.3.1 Methods

As detailed in Section 2.1.3, recent work around human interaction with non-specialist autonomous agents has adopted an in-the-wild approach to achieve realistic results (Yang and Newman, 2013; Alan et al., 2014; Bourgeois et al., 2014; Costanza et al., 2014; Yang et al., 2014). We agree with such an approach and we believe that it is important to run in-the-wild studies – indeed we do so in Chapter 3. However, as part of our work we are interested in exploring the opportunity to study interaction with autonomous agents in the lab. Lab studies allow for precise measurements to compare alternative experimental conditions, such as different interface features. Moreover, they tend to be faster and cheaper to run than field trials. However, studying user adoption of autonomous agents in a lab setting involves the challenge of maintaining a high level of ecological validity; to address it, we turned to experimental methods used in behavioural economics.

Typically, experiments by behavioural economists incorporate money in so called *choice situations*, i.e. situations where study participants have to make a decision out of multiple choices. For example, subjects would be asked whether to choose between an 85% chance to win $1000 (with a 15% chance to win nothing) and the alternative of receiving $800 for sure[2] (Kahneman and Tversky, 1984). More specifically our work was motivated by behavioural economics studies which involve the performance of repeated tasks with actual financial incentives: money is handed to participants based on their actions in the study. For example to investigate the effect of the perceived meaning of tasks on people's motivation to work, Ariely et al. (2008) designed a study where participants were paid to complete a simple task: assembling a Lego model. Participants had the option to repeatedly complete this very same task several times, but each time at a reduced wage rate (first $3.00, then $2.70, then $2.40 and so on) [3]. Ours is not the first HCI research to turn to behavioural economics for inspiration. Previously HCI researchers have suggested employing persuasion techniques based on effects studied by behavioural economists to promote healthy snack eating (Lee et al., 2011). However, our approach is different and novel in that we turn to behavioural economics for experimental methodology.

In more detail, we incorporated choice situations in the design of our studies described in Chapters 4 and 5. In these studies, participants have the choice to perform two tasks: one that involves the use of an autonomous agent, and the other without, both of which are paid. In such context, if participants perform the task involving the agent, they give up the option to perform the other task, with an associated opportunity cost, because they have a limited amount of time and tasks allowed in the study. Our aim is

[2]Even though the first choice has the higher potential gain, most participants would prefer the guaranteed choice. This was posed as a hypothetical question, no money was handed to participants.

[3]Their results show that manipulation on the task meaning would induce people to work for a significant lower pay rate.

to mimic a real-world situation whereby if a user chooses to invest time interacting with an autonomous agent, doing so would cost the user time and effort.

2.3.2 The Ultimatum Game

Going beyond inspiring our study methods, we also explored behavioural economics literature to inform us about how people deal with uncertainty as a result group interactions. This is particularly important for Chapter 4, where we are interested in applying autonomous agents in group buying situations. Since group buying situations involve some form of bargaining, it is important to understand how people typically act in bargaining scenarios. In the literature of behavioural economics (or more precisely behavioural game theory), there is a well-known experiment that has looked into this, known as the *ultimatum game*. The ultimatum game is an economics experiment that tested the importance of fairness in bargaining and ultimatum scenarios (Guth et al., 1982). The game involves two players, a *proposer* and a *responder*. The proposer is given an initial endowment (e.g. £10) and has to decide how much to split the money with the responder. The responder in turn then decides how to respond to the offered split: if the responder agrees, both of them receive the agreed split; if instead the offer is rejected, they both get zero. The rational action would then be for the proposer to offer the minimum amount allowed above zero, which should be accepted by the responder (as anything is better than nothing). The findings however suggest that people usually offer around 40-50% of the initial endowment and that offers below 20% are rejected half of the time (Roth et al., 1991).

Although there are many versions of the game (Camerer, 2003), the consensus is that fairness is important in bargaining or ultimatum situations. For example, Sanfey et al. (2003) compared how people responded to offers made by human partners against those of computer partners. Their findings suggest that while people accepted fair offers more than unfair ones regardless of their counterpart, participants accepted more unfair offers from computer partners than from human partners. However, in a different version of the game where participants had to encourage others to join their teams in fulfilling a collaborative task, it has been found that people tend to offer more fairly to humans than they do to agents (van Wissen et al., 2012). There are also versions of the game where it is framed in terms of losses (or savings), and studies have shown that proposers either offered similarly or more fairly than those who played the gains version of the game Camerer et al. (1993); Buchan et al. (2005). For instance, in the study by Neumann et al. (2017), proposers in the losses condition offered higher (45%) than those in the gains condition (40%), resulting in more agreements and higher payoffs for responders in the losses condition. Nonetheless, the offers in both conditions fall within the fair range, suggesting that regardless of the framing, people put importance on fairness.

Inspired by these findings, nechanisms that incorporate fairness have been developed to enable agents to effectively negotiate with people (e.g. (Gal et al., 2011; Chan and Chen, 2016)). For example, researchers have designed an agent that can adapt its behaviour based on the culture of its negotiation partner (Gal et al., 2011). There are also a number of mechanisms that take into account fairness amongst individuals in resource allocation problems. In particular, research has looked at how to achieve envy-free divisions of a divisible good amongst agents with different preferences (e.g. Roos and Rothe (2010); Cohler et al. (2011); Alijani et al. (2017)). For instance, Gal et al. (2016) developed and evaluated an algorithm that can determine the fairest and most envy-free solution to assign rooms to several house mates and divide the rent between them. While these works show that it is possible to achieve fair mechanisms, they do not touch upon the crucial issue of whether fairness impacts one's inclination to make use of agents acting on behalf of their human owners.

There are also some studies that examined how users interact with agents in an ultimatum-game setting, to examine whether agents can also encourage fairness amongst users. For example, in a study by de Melo et al. (2016), participants took the role of the proposer in a version of the ultimatum game. Some participants made offers directly to other human counterparts, whereas others made their offers via a programmable agent. Their findings suggest that people tend to make fairer offers when using an agent than when interacting directly with others. Similar results were also found in a more complex multi-issue negotiation task, where participants played the role of a mobile phone seller and completed negotiations either through an agent or directly to another person (de Melo et al., 2017b). The bargaining scenarios in these studies have been carried out as an individual task. However, in real life, bargaining scenarios often take place along side other activities. In our work, particularly as depicted in Chapter 4, we aim to develop a study method that would involve the bargaining scenario in a more realistic situation.

2.4 Summary

In this chapter, we covered various studies of how people interact with autonomous agents. In particular, we covered studies of interaction with specialist autonomous agents, agent-based decision support systems and autonomous agents in everyday situations. Our exploration of research around interaction with autonomous agents has shown the importance of considering how these technologies can affect domestic practices. However, as we identified, most of these studies have focused on applications such as doing the laundry, energy systems and cleaning robots. Moreover, most of these studies focused on research prototypes and little has looked at people's practices around already existing autonomous agents. Our review of HCI studies on domestic food practices has shown that research has investigated how technology can be used to support people's food practices, but not in terms of autonomous agents that do work (e.g. food shopping) on people's behalf. In particular, no research so far has looked at the factors that enables agency delegation, especially around food-based practices. To address these research gaps, we examined people's interaction of autonomous agents supporting their food practices. In particular, we focus on food shopping through veg box schemes (Chapter 3) and group buying situations (Chapter 4).

While research also suggests that people are inclined to adopt and delegate agency to autonomous agents, be it in specialist or non-specialist situations, little research has been able to address the issue of users over relying and not relying enough on autonomous agents. Prior work has suggested that displaying confidence information of systems can address such issue. However, most of these studies focus on context-aware systems rather than autonomous agents. Furthermore, these studies offer little ecological validity and fail to qualitatively understand how people employ confidence information to interact with autonomous agents. In our work, particularly in Chapter 5, we make use of displaying confidence information with the aim to help users adopt and correctly rely on autonomous agents. In addition, based on our literature review, we borrowed methods from behavioural economics to design controlled studies with high ecological validity, which are presented in Chapters 4 and 5.

Chapter 3

Observing Agency Delegation in a Domestic Everyday Setting

In this chapter, we present our work around understanding how people currently delegate agency to an inherently unpredictable autonomous service. In particular, we describe a qualitative study of everyday practices around an established service that exemplifies agency delegation, through the *veg box scheme*, in an effort to understand how the practices surrounding agency delegation in an everyday setting. This work is based on addressing R1 of our research questions: *How do people deal with uncertainty, particularly when it comes to delegating agency in everyday practices?* It is important and necessary to understand agency delegation, especially the factors that facilitate and hinder people into entrusting responsibilities to others. This is because such factors may allow us to formulate design implications for future autonomous agents, so that people will be willing to delegate agency to them, leading to the adoption of the technology.

In this chapter, we start by describing the motivation for this work, followed by the full description of the study. The findings and discussions are then provided, which is then succeeded by the design implications. Finally, we summarise the whole chapter and its contributions.

3.1 Motivation

Our literature review has highlighted the opportunities and the importance of considering how autonomous technologies can affect our everyday lives. In particular, these studies have focused on applications such as doing the laundry (Bourgeois et al., 2014; Costanza et al., 2014), energy systems (Alan et al., 2016a,b; Yang and Newman, 2013; Yang et al., 2014, 2016) and cleaning robots (Forlizzi and DiSalvo, 2006; Sung et al., 2007). With a few exceptions that have studied off-the-shelf products such as

Figure 3.1: A participant's photo report of a large veg box with fruits.

the Roomba (Sung et al., 2007), and the Nest Thermostat (Yang and Newman, 2013; Yang et al., 2014), prior studies around autonomous agents in the home have focussed on research prototypes. To complement this prior work and to inform the design of future autonomous agents we study everyday practices around an established service that exemplifies agency delegation, particularly through the veg box scheme, which we described in Section 2.2.2.

In subscribing to the veg box scheme users delegate to the service provider the decision on which food items they receive. We therefore see this as an instance of agency delegation similar to other existing autonomous services, such as a thermostat to which one might delegate the decision to heat or cool a home. Moreover, we focus this particular work on agency delegation in food-based practices, and we see the veg box scheme as a relatable autonomous service that already exists in the market and also has an element of unpredictability. Studying people's use of veg box schemes allows us to observe practices around autonomous services that have been adjusted over time, enabling us to investigate how people deal with such unpredictable services in everyday life. In the next section, we specify the details of the study.

3.2 Study

For the purpose of our study, we worked with a community farm running an organic veg box scheme in the UK. The farm helped us to recruit participants from their existing customer base by advertising our study on their marketing email list and through their social media accounts. The veg box scheme run by the farm offers a relatively limited range of options (see details in Section 3.4). Some of the produce in the box are grown at the farm, some at nearby farms, while some are from abroad, with different proportions week by week. We felt that such limited options fitted well our intent to explore a situation where the veg box service has potential to perturbate existing food practices, allowing us to investigate the following: how do people integrate autonomous food services to their existing food practices and how do people make the best use of the veg box? Addressing these will allow us to formulate implications for designing future food-related autonomous agents and technologies to support services such as the veg box scheme. The study has been approved by the Southampton Ethics Committee (ref: 19339).

3.2.1 Groups

Our participants were divided into two groups.

- *Subscribed* – This group included five participants who were already subscribed to a veg box scheme. We recruited most of them from the community farm we were working with for the duration of the study. One participant was subscribed to a different company, but her subscription was similar to the veg box scheme offered by the partner farm.
- *Non-subscribed* – The *non-subscribed* group instead contained the other six participants who were interested in veg boxes but were not subscribed to any at the time. Through the social media accounts, mailing list of the farm and word-of-mouth (via existing participants). We subscribed these participants to a veg box scheme from the same farm (for consistency) based on their household size, and their preferences, such as whether to include fruits and potatoes.

The purpose of the two groups is to enable us to observe both established food practices around the veg box (from the *subscribed* group), but also emerging ones (*non-subscribed* group).

3.2.2 Method

Data collection involved the following three phases:

Entry Interview. We firstly scheduled an entry interview with each participant. In this interview, participants in the *subscribed* group were asked questions about their motivations for subscribing, how the veg box affects their food practices and also other opinions related to the service. Participants in the *non-subscribed* group were instead asked about their existing food practices and their knowledge around veg boxes in general. After the interview, they were given instructions regarding the diary study phase. Each interview lasted between 14 and 36 minutes.

Diary Study Period. During this period, participants were asked to report for 14 consecutive days when they used or disposed of any veg box content, as well as any non-use, such as eating out or making food that did not involve the veg box. The information was reported through Whatsapp[1], a popular chat application freely available on most mobile platform. In fact, all but one of our participants were already Whatsapp users. To send a diary report, participants had to take a photo of the ingredients or the dish they made and annotate the photo (either through text or audio) detailing a short description of the photo (e.g. Left of Figure 3.2). Participants were also given the option of sending a video as an alternative. We offered participants to send daily reminders, and 7 of them accepted. The choice of the diary study allowed us to observe (to some degree) participants' actual use of the veg box for some period of time.

Exit Interview. Finally, we scheduled exit interviews for all our participants. In particular, we were interested in understanding how our participants went about the decision to prepare the dishes involving the veg box contents. Participants in the *non-subscribed* group were also asked about their thoughts of the veg box and the likeliness of them continuing the subscription. Each interview lasted between 18 and 56 minutes (on average 30 minutes, SD: 10).

Although participants only received 2 boxes during the diary study, participants were asked to send reports when an activity relating to the veg box occurs. Since, we expect that such an activity would happen multiple times in a day, we opted for a 2-week long diary study. In addition, our study does not only rely on the diary study, but also on two interviews. Indeed, other studies that utilised such an approach where the reports are also relatively frequent also opted for a duration choice that is not considered to be long (e.g. Cecchinato et al. (2016)).

3.2.3 Participants

We recruited a total of 11 participants (10 females, 1 male) through word of mouth and with the help of the farm (as described above). Our study covered a variety of household sizes as detailed in Table 3.1. Furthermore, each of our participants considered themselves as the person mainly responsible for all food-related tasks in the household.

3.2.4 Reward

Participants received up to £60, according to a reward scheme based on their group and their level of engagement in the study. Those in *subscribed* group were paid £32 for taking part in the study, plus an extra £2 for each day that they report at least once. Such a way of incentivising participants to motivate them to send reports was inspired

[1] https://www.whatsapp.com/

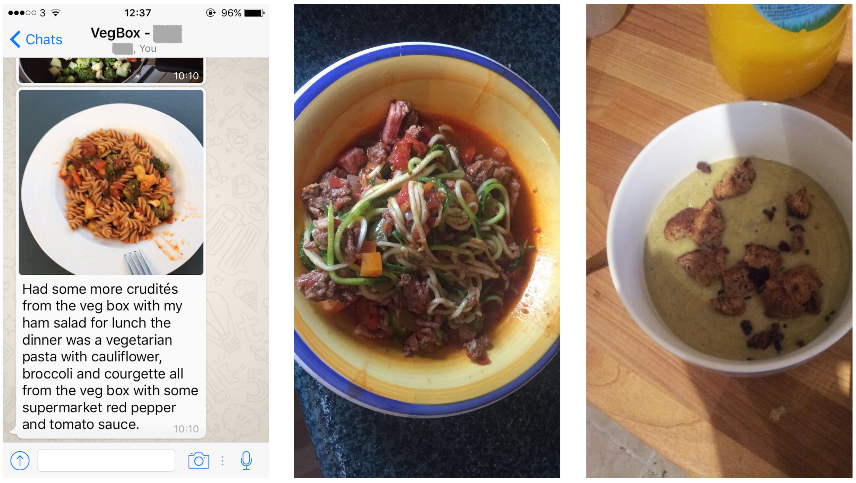

Figure 3.2: (Left) Screenshot of a participant's diary entry (Ella). (Middle) "Courgetti" – The spaghetti noodles were replaced with spiralised courgettes (Daphne). (Right) A photo report of a soup made to hide the taste of some vegetables (Emily).

Table 3.1: Table of participants. Others indicate the number of household members who are regularly involved in the participant's food practices. One participant, Daphne, switched their box size from small to large (because of the kids not being at home the first week).

	Participant	**Age**	**Others**	**Size**	**Freq.**
Subscribed	Ella	40s	1	small	weekly
	Tanya	30s	2	small	fortnightly
	Abby	30s	2	medium	weekly
	Sofia	30s	4	medium	weekly
	Emily	20s	1	medium	weekly
Non-subscribed	Ines	20s	1	medium	weekly
	Laura	30s	-	small	weekly
	Andrea	30s	3	large	weekly
	Daphne	40s	4	large	weekly
	Liam	40s	-	small	weekly
	Zoe	20s	-	small	weekly

by a previous diary study (Sohn et al., 2008). Participants in the *non-subscribed* group were provided with 2 weeks worth of free veg boxes, roughly costing £20 per veg box. Additionally, they received an extra £1.50 for each day that they report at least once. All cash incentives were handed to the participants during the exit interviews.

3.3 Diary reports

Overall, 7 participants were able to send daily reports for 14 days. The remaining 4 participants missed at least one day of reporting, of these, two told us that they forgot to report, one had to be in hospital, whilst the remaining one cancelled her subscription

Table 3.2: Summary of diary reports across the entire study. Note that the reports were categorised based on their content.

Categories	Frequency		
	Sub.	Non-sub.	**Total**
Meals involving veg box items	72	71	**143**
Veg box items thrown away or given	12	3	**15**
Veg box items eaten raw as snack	6	41	**47**
Veg box items difficult to identify	0	2	**2**
Reports of non-use (e.g. eating out)	29	19	**48**

towards the end of the study (details below). Participants sent 29 reports on average over the course of the study. In more detail, participants who were reminded about the diary study sent about 30 reports on average over the course of the study while those who didn't ask for a reminder sent 28 reports on average. Because the difference between the two is minimal, we believe that no self-selection bias was introduced by letting users select whether they were notified or not. It is also worth emphasising that the reports from the diary study were used primarily as prompts for the exit interviews, rather than as primary data. However, to give a sense of the variety of information that we received, and the level of engagement of our participant with the study and with using the box content, we provide a summary in Table 3.2.

3.4 The veg box scheme

The farm offers a number of options for the veg boxes, ranging from small (6-7 items) to large family size (9-10 items) veg boxes (Figure 3.1). Each veg box is available in two variants: with or without potatoes. Fruit boxes can be included as part of the veg box or as separate. The farm provides a list of upcoming items 2-3 days before the veg box is delivered via email and on the website. Customers can 'blacklist' produce they dislike, which would then be replaced, however, this option is not publicly advertised, and it is made available only on a case-by-case fashion. Although customers can choose the size and type of the box and whether to include or exclude certain produce, its content is still decided by the farm. Users can buy additional items to top up the box, but not to replace the boxes' contents. The farm also provides various options for delivery frequency, ranging from once a week to once a month.

Participants in the subscribed group learnt about the scheme through local promotional events, word of mouth, or via online searches. Table 1 shows the subscription choices of our participants regarding the size of their veg boxes and delivery frequency. Participants in the subscribed group accounted for how they came up with their subscription choices based on two factors: 1) their usual vegetable consumption, and 2) their anticipation of how long produce will stay fresh. We made subscription choices for the non- subscribed group based on their preferences to include or exclude fruits and potatoes.

The farm offers two options to receive the boxes: either through home delivery for a small fee or for free from a number of collection points. Collection points include local churches, cafes, shops and the farm itself. As noted above, the farm introduced a small fee during the study (less than the delivery fee) for collecting the veg box from anywhere other than the farm. For convenience and as part of the compensation for their time, we decided that our participants in the non-subscribed group would receive the boxes through home delivery.

Several members of the non-subscribed group gave additional instructions about receiving the box, such as leaving them with their neighbours or making sure that the box was not left in a completely visible location to avoid theft. It is also worth noting that some of them expressed willingness to collect the boxes from a collection point. In contrast, all but one of the participants in the subscribed group collected their veg boxes. This delivery choice was influenced by monetary cost, although participants reported that collecting the boxes was inconvenient. Produce in the veg box is largely grown at the farm or at other farms in the locality, though some (particularly fruit) comes from further afield and even abroad.

3.5 Findings and discussion

Herein we present the themes that emerged from an inductive thematic analysis of the data (Braun and Clarke, 2006). Two researchers were involved in conducting the interviews. The interviews were audio-recorded, fully transcribed. The transcripts were then open coded and the open codes were grouped in broader categories into key themes that *'make or break'* agency delegation in everyday life that are presented in the following.

3.5.1 Exercising some control

Clearly the opportunity exists for users of the veg box scheme to exercise some control over its inherent unpredictability by a) expressing preferences in the course of box selection, and b) blacklisting items. For example, Zoe provided an example of how they expressed their preferences in the course of box selection.

> Zoe: *"Yes, in general and week by week, so I had quite a lot of potatoes then I'd be like actually I've got potatoes, I don't need any more so don't send me any this week, send me something new."*

The statement accounts how exercising control over items in the veg box can be occasioned. Here, already having items available that might be included in the veg box occasioned explicitly flagging potatoes as undesired in that week. Yet, it is clear that the control is limited to declaring undesired items; the participant orients to the remaining unpredictability in what they might receive instead (*"send me something new"*).

Only one person took advantage of the option to blacklist items.

> Tanya: *"They did once give us this, like, really awful green tomato. I'll eat anything that's fresh, but they were awful. I let them know on my account. I was, like, don't give me them ever again, please."*

3.5.2 Warranting delegation

It would thus appear that to participate in a veg box scheme is to accept unpredictability as a condition of engagement, but it is important to appreciate that in doing so agency delegation is *warranted* on various grounds, i.e., participation in agency delegation is subject to various conditions being met. Here, we orient to these warrants made available in participants' accounts.

Firstly, our participants warranted agency delegation on grounds of "locality". Tanya: I quite like the idea of the [anonymous] community farm because it's about, you know, local people growing veg. And I like the whole ethical ethos behind it. This concern with the local as a warrant for agency delegation glossed a number of interrelated drivers, such as the motivation to support "local people" and the connected "ethical ethos". Relatedly, participants explicitly oriented to a concern with seasonality of the produce.

> Daphne: *"I like the idea of having that seasonal approach of what you've got and you have to make your ... because when you shop at a supermarket all the time everything's always available and so you don't shop, you don't cook seasonally, which is a shame. And so I quite like the fact that you're getting stuff that is obviously in season as far as possible, so it gives you an idea of what actually is around (...)."*

Daphne here expresses their appreciation of seasonality as a value provided by the veg box distinct from the *"everything's always available"* value provided by supermarkets. Thus, agency delegation was warranted in terms of buying local produce, and high quality and seasonal local produce at that, and was something that therefore supported the local non-profit farm, as well as the local community more broadly.

Other participants also warranted agency delegation on grounds of enabling a *"healthy diet"*. For example, one of our participants had recently turned vegan, and took advantage of the veg box as an opportunity to enforce her diet:

> Emily: *"That is kind of part of the reason why I started getting the box, because when I went vegan I wanted to make sure that I had a varied diet, and I thought this would like force me to have some vegetables I would never buy because maybe I don't like them that much but I'll still eat them."*

Others warranted agency delegation on the more mundane grounds that the veg box scheme provided them with "peace of mind", in ensuring that they always had a supply of fresh fruit and vegetables to hand.

Agency delegation was also warranted on the grounds of "value for money".

> Ines: *"I think it's fair enough for the price. The products are very good, the taste is awesome. It's not the same taste when you buy it at the supermarket, it's not. You can feel the taste, the real taste, so I think it's very, very fair enough. And I don't think it's too expensive - I think it's very cheap, because they bring it to your home, you don't need to go out. You don't need to carry the bags. They take your job."*

As Ines makes visible, value for money is not simply reducible to matters of financial cost, but includes the other costs implicated in going and getting food for oneself (particularly time and labour). Ines also highlights the taste of locally grown produce, which was a topic that was brought up by others as well.

> Andrea: *"I mean getting to know the boxes over two weeks I could see what I would really like to have and it's like the cucumbers maybe and the marrows and those things that grow so well here and they actually taste so much better just being grown right here."*

Not only does Andrea appreciate the taste of the locally grown veg, they also orient to discovering the things they really liked. These elements of discovery and serendipity frequently came up. Thus, one of the key drivers of agency delegation was the "element of surprise" occasioned by the veg box.

> Sofia: *"The box exposes us to vegetables that otherwise we wouldn't try or think about buying or even have available in the supermarket, like the Swiss chard or the rhubarb or kohlrabi. Like I had never tried it before, right, because I'm not from here and the vegetables are different, I would just never end up trying them. So that's kind of good, that you are forced to actually use the things you don't know."*

Our participants broadly welcomed the unpredictability of the veg box and how it enabled the discovery of new food items.

> Liam: *"I guess it just makes you do things differently than you would do. Yes and I like most vegetables. Just means you have a variety that's different things. You wouldn't necessarily choose to pick up yourself. I sort of like that about randomness aspect of it."*

Liam here highlights that he liked the "randomness aspect" of the veg box which occasioned "doing things differently".

Related to this, the effect the veg box had on participants' cooking routines was valued explicitly.

> Daphne: *"I might well consider restarting because it was quite, I did quite like the challenge of using up and also because I had got my milk … because I stopped getting my food being delivered because I'd got into a rut and so I quite liked, and so I suppose it's quite nice to have a sort of jolt out of your rut. And to have to sort of consider more what you're, put a little bit more thought … because again, I like cooking so it's quite nice to have that sort of challenge."*

Daphne here describes appreciating the veg box in that it challenges the usual cooking routine by providing a *"jolt out of your rut"*, which encourages to consideration and thought.

Thus far, participants' accounts have made the grounds upon which agency delegation turns visible, including valuing the local, seasonal ethos the box brings, which people have associated with a healthy diet, peace of mind, value for money, and taste. In particular, 'the element of surprise' and discovery of new items, as well as the challenge this provided for people's routines warranted accepting the uncertainty of agency delegation. The next theme explores how our participants got on with incorporating the veg box in their everyday life.

3.5.3 Incorporating the veg box in everyday life

It was clearly the case that the unpredictable has to be *incorporated* into the orderliness of everyday life and the activities that constitute it.

> Tanya: *"During our weekly shop I will have a look at the menu, and the list of vegetables - because they release that, I think, on a Tuesday - and I'll just*

> *see what we're getting and then what else we need. Dependent on that, we'll buy whatever else I think that we want to have that week."*

The "menu" Tanya speaks is her meal plan for the week ahead. Of course, not everyone makes meals plans, but all of our participants routinely complemented the veg box to varying degrees to enable its incorporation into everyday life. Our participants commonly bought staple items that they did not receive in the box to permit a range of recipes, and many (like Tanya) bought specific items to match what was received in the veg box to provide for specific meals. The veg box does not stand-alone then, but is *planned in* to grocery shopping and frequently *built* in to the delivery of specific meals. The unpredictable is made at home in an orderly world through the intentional weaving of veg box contents with other items to meet prospective need. Thus, agency delegation turns on the ability for the unpredictable to be made *accountable* to the particular demands and needs of everyday life: e.g., what we want to eat next week, and even what we want to eat on particular days or at particular times.

It may be the case that the items that get delivered are things the participants would routinely buy and that the veg box fits in with participants' usual meals. However, it may also be, and indeed often was, the case that situations arise where participants have to *adapt* recipes in order to fit the contents of the veg box in with their meals.

> Emily: *"So whenever there is something that comes in that I really don't like the taste of I just put it in the soup, because then you can hide the taste (Right of Figure 3.2)."*

Participants routinely adapted recipes, and for variety of reasons both positive and negative, including people not knowing what to do with particular vegetables, not liking them, getting children to eat them, and using up leftovers. The need to adapt what's in the box also fosters innovation.

> Daphne: *"As I say, we had the courgettes and my husband's not a massive fan of courgettes until I tried spiralising them. Then he decided actually he loves those spiralised courgettes and, well, we can now replace spaghetti with courgettes spiralised. So that was great. That was a good discovery to make."*

Daphne's "spiraliser" is a kitchen gadget that allows her to cut vegetables into long pasta-like ribbons. Whether spiralising courgettes (Middle of Figure 3.2), or turning nasturtium flowers into spicy pakora, or making rich tomato-based vegetable sauces, etc., our participants innovate in order to incorporate the unpredictable into their usual practices, sometimes (but by no means always) resulting in "good discoveries". In describing such "good discoveries" participants at times explicitly oriented to the sensual delight of recipes improved with items from the veg box.

> Liam: *"I often have pasta just with tomato sauce and tuna, which as now I will throw lots of vegetables into it, which will significantly change the amount of vegetables I have in my diet. That was a very easy dish and I was surprised at how nice it was to be crunching on vegetables in a pasta dish."*

Liam here describes the surprise and sensual delight they experienced from augmenting a simple dish they routinely cook with veg box items. Thus, incorporating the unpredictable into the normal routine *requires* innovation as it seems to us that the unpredictable *has* to lend itself to "good discoveries" if it is to be sustainable. Simply compelling people to adapt to what they are given - e.g., through "hiding" unloved vegetables in soups is likely to result in agency delegation being revoked.

A further theme that emerged from the study speaks to the barriers that encumber incorporating agency delegation in everyday life.

3.5.4 Barriers to agency delegation

There is a strong sense then in which agency delegation is constrained, not only in that an agent's actions must be accountable to people's existing routines but also in that its actions must facilitate such routines, and it is in the latter respect that the veg box becomes problematic particularly amongst the non-subscribed group. These constraints may, then, be seen as barriers to agency delegation or organisational features of it that must be accommodated on the pathway to adoption. Whether they are accepted or rejected turns on their alignment with the grounds that warrant agency delegation in the first instance.

> Andrea: *"They bulk it out with some of the fruit, obviously stuff that they've sourced elsewhere because it's - you wouldn't grow bananas and melons locally. The bananas are from Peru, so same as buying from anywhere really."*

Andrea was not the only participant who noticed that fruits in the box were often not sourced locally, nor who thought it worth the extra cost and commitment when you can buy such things "anywhere", and for less.

> Tanya: *"I don't actually think it's cost effective, particularly sometimes for the quality of the food that you get and the amount that you waste. Because, you know, the Co-op [a supermarket] does three items every week where it charges, like, 69p. So you can get a bag of potatoes for 69p. You can get a courgette, or a broccoli, for 69p. Tomatoes can be 69p. That's nothing for vegetables, and actually the quality's very good."*

Agency delegation must not only comply with the grounds upon which it is warranted then, be it a concern with locality or value for money (etc.), but insofar as financial control is being delegated then agents must also demonstrate that they are "cost effective". As Tanya makes visible, this goes beyond a concern with money alone. It is money in relation to other calculable benefits that matters, e.g., that much the same quality can be had and with less waste by shopping at the Co-op.

The issue of waste was a major concern for participants in both groups.

> Liam: *"I'm often in a rush, so I figure I probably have less food waste if I buy pre-made things from Marks and Spencer's [another supermarket] and take the hit on extra packaging than if I was to buy lots of fresh food to cook and then waste what I cook and things."*

For Liam, the issue of waste comes down to a *calculable trade-off* between creating food waste or packaging waste. Most of our participants, like Liam, expressed concern about getting more than they might use in a veg box, and when this occurred some took active steps to avoid food waste.

> Zoe: *"I just put them out on the table. [My housemates] helped themselves, and whatever wasn't taken then they were put in the compost."*

Five of our participants gave, shared or swapped veg box contents with housemates, neighbours, friends or family to avoid food waste. Four also composted spoiled or undesirable food items in a bid to recycle them. For the rest, some veg box content inevitably ended up in the bin.

> Sofia: *"We have situations that's like, what do we do with them? It's like, I don't know, let's not do them today. So I think they probably - they might go to the garbage, just because we are ignorant, or lazy to Google it."*

It is also the case that Liam was not alone in often being *"in a rush"*, and this too is consequential to agency delegation.

> Andrea: *"It's a nice idea but modern life gets in the way, and that's the key thing with these veg boxes I think. It's a bit of - I kind of want to know what I'm going to get in a way. Yes, I want to be inspired - there's these other ideas which my brother's actually subscribed to which is "Hello Fresh", they send you a recipe as well and all the ingredients to make up that meal. So that's the other thing. It's like, yes, you get a random thing, and you want to*

> *be inspired to make something new, but then you have to think that you have to go and buy your other things to go with it. So that's going to be tricky sometimes."*

As Andrea makes clear, calculation also includes the potential for agency delegation to gear in with her *"modern life"*. Incorporating the veg box into everyday life then, can be burdensome and makes visible how agency delegation can be a double edged sword; it's not always straightforward to innovate or adapt meals around veg box items, it can be associated with additional hassle and cost (*"having to go and buy your other things to go with it"*). Others voiced a related concern more succinctly in terms of *timing*.

> Daphne: *"My only problem with it was is it arrived on a day... it, like, arrived in the... towards the end of the week, and I'd already done my shopping at the beginning of the week, so sometimes I'd end up with... so that... so I think I found it slightly awkward in terms of that."*

The timing between their shopping and the arrival of the box has made things *"awkward"* for Daphne. All of the participants who were not regular subscribers commented how they would find it difficult to commit to a regular subscription because of misalignment of veg box delivery timing and their schedule. It is not only a matter of having to go and buy other things then, but of having the opportunity.

Agency delegation thus turns in significant respects on its ability to gear in with people's busy schedules and this latter point speaks to the observable embeddedness of agency delegation within a social division of labour. Thus agency delegation turns on its *coordination* with *multiple actors* and their *activities*. Whether it be gearing into schedules, or shopping opportunities, or just who is doing the cooking today and for which people, or even eating out, our participants made it perspicuous that agency delegation is accountable to and must mesh in with the social division of labour in which it is embedded and resides.

3.6 Implications

In this section, we present a number of design implications for future autonomous agents. In particular, the findings were discussed with other project members in order to come up with implications that are novel and relevant to supporting the adoption and delegation of users to autonomous agents.

Based on our findings, we can see that agency delegation is *warranted*, i.e., that it turns upon some justification for surrendering autonomy and delegating agency to another party. We can also see that it must be possible to *incorporate* agency delegation into existing everyday activities. And we can see that agency delegation is subject to *constraints* that can act as *barriers* to adoption. The ways in which agency delegation is warranted and incorporated into everyday life are therefore accountable to a range of calculable concerns that determine its fate.

3.6.1 Supporting creativity

While automation is typically driven by values of efficiency and convenience (Parasuraman and Manzey, 2010; Costanza et al., 2014), our data revealed that a different set of values and benefits could be derived from delegating agency to an external service. Many of our participants found the "element of surprise" of the unpredictability of which vegetables their box would contain made them more creative in their cooking. Others valued the discovery of previously unknown vegetables. Others were delighted by the sensual experience of enhanced recipes. This suggests that a certain degree of unpredictability in agency delegation can speak to a different set of values: delight, discovery, and creativity. This result points to opportunities for autonomous technology to offer users items or actions that they may not expect as a feature. It also suggests that users of autonomous agents may be tolerant to unpredictability, at least if its cause is understood and accepted (in the case of the veg box, the seasonal availability of items).

We feel encouraged then to stress that the design of autonomous agents should not only strive towards optimality, accuracy, etc., but also accept that uncertainty can in fact lead to beneficial user experience. This finding chimes with Gaver's work on ambiguity as a resource in design (Gaver et al., 2003), as well as Roger's appeal to move Ubicomp applications beyond the agenda of calm computing towards applications that foster creativity (Rogers, 2006). There is much room then for new agendas in domains typically motivated by convenience and efficiency, such as smart homes and the Internet of Things (IoT). Work in this space has already begun, for example to explore how the unpredictable can delight and surprise people, and how it might support creative practice (e.g. Mennicken et al. (2016)).

So what can design do? Insofar as unpredictability is welcomed into everyday life, then we would also add that supporting the practices of *adaptation* and *innovation* is an important ingredient to add to the design mix. Users will need to adapt and innovate to be able to incorporate the unpredictable into their everyday life. Without them, and the serendipitous discoveries that often accompany them, there is no solid ground for agency delegation to stand on and the warrant will inevitably be revoked. After all, people do not just want to make do with what they are given, especially food. And how might we support adaptation and innovation through design? Considering the food domain, creative ideas how to prepare, combine, and cook items could be provided[2]. Aside from

[2] An example is https://www.ibmchefwatson.com/

recipes, this could include instructions how to adapt existing recipes, how to manage a surplus of items to avoid waste, or how you might ‘hide’ flavours of undesired items.

3.6.2 Supporting value calculation

More in general, in a domestic food context, our study has shown that locality, seasonality, health, value for money, serendipity, and creativity all warrant agency delegation with respect to food. The grounds are not static either, but subject *continuously* to calculable concerns to do with warrant compliance, and cost effectiveness with particular respect to quality and waste. It might be said, then, that a heterogeneous *array of values* motivates and sustains or curtails agency delegation. The array is important. No particular value or combination of values is ubiquitous. It all depends on people’s habits, on just what these people, in this house, value. There is need then to cater for all and support people’s value calculation practices through design to enable effective agency delegation, where effective means that agency delegation can effectively be incorporated into the existing practices. Rather than to attempt creating autonomous agents that are assumed to act flawlessly, we advocate for mixed-initiative approaches (Alan et al., 2016a), in which agency can be fluidly transferred from the system to the user and vice versa to support the user’s own value calculation.

So how would or could you enable value calculation through design? One approach could involve solutions that allow people to express their values, and define how computational agents should respond to them. Technical approaches that have been explored in the literature include value sensitive design in agent design (Friedman et al., 2008; Friedman, 2013), and preference learning. For example, PlateClick employs a quiz-based user interface to elicit people’s food preferences, which can then be used in an online learning framework (Yang et al., 2015a).

An alternative approach is to make information relevant to the values and to the device or service easily available to the users. For example, in the case of autonomous agents that mediate domestic *energy consumption*, data about financial savings on the utility bills, about fuel efficiency and about the origin of the energy (e.g. gas, wind, solar) could help users to continuously evaluate the system’s performance and hence grant or revoke agency. Such an approach has recently been demonstrated by a small variety of research projects (Alan et al., 2016a,b; Bourgeois et al., 2014; Yang et al., 2016). Similarly, for autonomous cars, it may be useful to display information related to the financial cost of each trip (such as in Southern et al. (2017)), the fuel efficiency, and how long the trip will take - for all these variables the comparison between autonomous operation and manual mode should be shown.

More in general, we see potential for autonomous agents to integrate techniques that can track the *provenance* of goods and the derivation of decisions, such as distributed ledgers (Adviser, 2016). For example, in the case of the veg box, our data suggests that customers question and make assumptions about where the produce comes from[3] (Moreau and Groth, 2013), and why it is included in a delivery. Is it local, or does it come from far away? Does it come from an environmental and socially sustainable supplier? A recent project considering this approach is Bitbarista (Pschetz et al., 2017), which explored the application of cryptocurrencies (based on distributed ledgers) to link the users of a coffee machine to remote farmers growing the beans.

3.6.3 Supporting local coordination

Our findings highlight the challenges of actively integrating agency delegation into everyday activities. Hence it is not sufficient for a computational agent to 'merely' attend to its delegated business: that business must also be made accountable to other mundane matters implicated, in this case, in the provisioning and consumption of food items. Thus, agency delegation around the veg box must integrate with household shopping, to ensure that meals can be provided, and even with specific plans as to what will be eaten and when. Again, incorporating agency delegation into everyday life is an ongoing matter conducted in response to calculable concerns with scheduling and coordination. Thus, integration extends to gearing agency delegation in with the social division of labour.

There is need, then, for design to support *coordination* to enable effective agency delegation, where effective here means that computational agents can demonstrably mesh their actions in with other local actors and (food related) activities. Research has demonstrated the design of agents that draw on further digital resources to enrich the grounds upon which agents take action. Resources have included for example people's calendars as a way to coordinate availability (Neustaedter et al., 2009), and people's location as a way to provide services just-in-time (Sohn et al., 2005), invoking the trope of "remember the milk" services that remind people at just the right time when they are near the shop on their way home (along the lines of e.g. Horvitz and Krumm (2012)).

The IoT offers potential to extend this work, taking into account the operation of (Internet connected) domestic appliances and smart jars able to sense and report their own content (e.g. Fan and Truong (2015)), as well as harvesting information from social media and shared calendar accounts, and from personal device location. This information could then be combined with data about available supply to extend stock control and supply chain optimisation through "the last mile" to the home, not only in terms of incoming goods, but also facilitating food sharing and food waste reduction, extending existing practices (Ganglbauer et al., 2014). It should be noted, however, that designing a system to integrate such heterogeneous sources of information in a timely and relevant

[3] Tracking the provenance of food has also been the aim of several companies, such as Project Provenance Ltd (https://www.provenance.org/)

manner is still a technical challenge. However, the point we wish to make is not that all solutions lie in automation, but to provide enhanced digital resources for people to make their own decisions, to provide resources that support people's own coordination practices.

3.7 Limitations

We do not suggest that this is all there is to agency delegation. In our study, agency is delegated to another human, not to software. As a result, direct comparisons between autonomous agents and veg box schemes cannot easily be made given that, unlike autonomous agents, veg box schemes are not based on sensor data collected from users. Other important features of autonomous agents, such as the ability to continuously monitor data and act autonomously based on those data, are missing in veg box schemes. Another limitation of our study is that it only covers a particular instance of autonomy. Other services, like the recipe boxes (detailed in Section 2.2.2), offer a higher level of autonomy, as users of recipe boxes have less control over their food choices in comparison with the veg boxes. Because of this, people's use of this kind of autonomous service may differ with how the veg box service was used by our participants. Users will likely tolerate "surprises" from a recipe box differently. Moreover, people's motivations to subscribe to the recipe box are likely dissimilar with the motivations of veg box users.

3.8 Summary

In this chapter, we reported the findings of a qualitative study, through semi-structured interviews and a two-week diary study with 11 households in the UK, which seeks to understand how people manage a veg box scheme as an instance of an inherently unpredictable service. In particular, we focused on how people manage agency delegation and integrate the veg box into their everyday life. Our findings suggest that agency delegation must be warranted, that it must be possible to incorporate delegated decisions into everyday activities, and that delegation is subject to constraint. We consider the potential impact of these social organisational issues on the design of future autonomous agents supporting food-based practices in the home, and the challenges of making agency delegation accountable to meal planning, persons' schedules, food-centred values, adaptation and innovation, and the social division of labour in which computational agency will ultimately be embedded.

Existing research has suggested that people are ready to integrate autonomous agents in their day-to-day routines, especially because of the increased efficiency and convenience that these agents are able to provide. However, this study extends prior work, as it further suggests that agency delegation can also be facilitated by factors that go beyond efficiency and convenience. In particular, our findings revealed that supporting user creativity and embedding human values in the design of autonomous agents can also encourage people to adopt agents and delegate agency to them, even when the actions of such agents are subjected to uncertainties. Indeed, this seems to suggest that an agent should not only focus on completing its goals for people to adopt and delegate agency to it, but it should also complete its goals with human values in mind. Therefore, it would be interesting to explore whether embedding human values in the design of autonomous agents impacts their adoption, especially in comparison to agents that focuses only on their goals. While the concept of designing agents that take into account human values is not entirely new, no empirical research to our knowledge has explored the impact of different agent behaviours in the adoption of the technology. In addition, as has been mentioned, an important limitation of this work is that it does not involve interacting with an actual autonomous agent. Such interactions will likely be different, given that people inherently treat computers differently than other people. This paves way to our next study, which we describe in the next chapter.

Chapter 4

Embedding Human Values in Agent-Design

The veg box study reported in Chapter 3 shed light on how people currently delegate agency to services that exemplifies autonomy in the context of food-based practices. One of the interesting implications of our previous study is that *human values* appear to affect one's inclination to engage people in agency delegation. In this chapter, we are interested to explore such implication, and understand whether designing an agent with human values affects a person's inclination to adopt it, especially in comparison to an agent that is designed to behave in a goal-oriented manner. In particular, our aim is to move to a more concrete perspective, through studying agency delegation by observing people's interaction with actual autonomous software agents. This study allows us to address R2 of our research questions as stated in Section 1.3, as it reveals the characteristics that a software agent should have for people to adopt it. In addition, the study method depicted in this work is dedicated for addressing R4.

The chapter begins by stating the motivation of this research. We then provide details of the study and also of the results of the evaluation. After these, we provide a discussion of the findings and formulate design implications for future shopping agents. Finally, we conclude this work by summarising the chapter and by relating its contributions to the literature.

4.1 Motivation

Similar to the veg box scheme, we are interested in agents that do work on behalf of the user, where such work has financial implications (e.g. the agent directly using the user's money to buy something). To inform the design of autonomous agents that make financial decisions on behalf of the user, we turn to research in Behavioural Game Theory

(BGT), a discipline that studies people's attitude towards making financial decisions and deal-making. For example, the Ultimatum Game (Guth et al., 1982), as described in Section 2.3.2, revealed that people's financial decisions take into account factors that go beyond rational choice, such as fairness and equality. Hence, in this chapter we focus on the specific issue of whether autonomous agents should be designed to mimic the "irrational" fair behaviour of people? Or would it be more *transparent* to design them as, perhaps simpler and predictable, machine-like goal-oriented machines? While previous work in the AI community has sought to develop various fair mechanisms (e.g., envy-free or egalitarian (Roos and Rothe, 2010; Cohler et al., 2011; Chan and Chen, 2016)), such work does not touch upon this crucial issue of how to design the interaction with agents acting on behalf of their human owners.

To address these questions, we focus on autonomous agents for grocery shopping, much like the veg box scheme we studied. To enable us to explore the concept of fairness in grocery shopping, we opted for a shopping scenario with a group setting involved. In particular, we considered the opportunity for autonomous agents to facilitate "group buying", or "collective buying": the collective purchasing of wholesale products (e.g. rice in packs of 25kg) to be shared across different buyers at a negotiated split (e.g. 1 kg per person). Group buying yields financial savings for buyers and has potential to lower financial and environmental cost for the suppliers (e.g. fewer deliveries). Group buying requires finding others interested in acquiring the same items, and negotiating with them how to share the goods and the cost. Such process can potentially be time-consuming and complex[1], so we see potential for autonomous agents to facilitate it.

Against this background, in this chapter we report on a comparison of different design choices for autonomous group buying agents. The study employed a combination of quantitative and qualitative methods, and, crucially, financial incentives related to the participants' performance in order to achieve higher ecological validity (an approach also borrowed from BGT). In particular, 20 participants were exposed to two alternative group buying autonomous shopping agents: one designed around fairness, and the other around being goal-oriented. To ensure some level of realism, we included multitasking in the study design. The results suggest that while people are likely to reject goal-oriented agents more than fair ones, people still welcome the use of a group buying agent as long as there are some guaranteed benefits. We detail the study in the next section.

4.2 Study

The study was designed to expose participants to two alternative group buying autonomous shopping agents, one designed around fairness, and the other around being

[1] As an example, group buying for energy schemes can take a couple of months before completion (https://goo.gl/CmLWPr).

goal-oriented. In particular, we wanted to examine whether these different behaviours can affect people's inclination to adopt and utilise the software agents. To explore this issue, we contrasted two alternative agent designs, which we further describe in the following subsection. This study has been approved by the Southampton Ethics Committee (ref: 31067).

4.2.1 Group buying agent

We designed a software agent to help users purchase items through group buying. When setting up a group buying agent, users were required to specify a number of parameters, such as the number of units they would like to receive and how long they want the agent to search (deadline). The role of the agent is to offer a price on behalf of the user to pay for a specific item. The agent then browses the market and looks for a group that it can join. Once it joins a group, it then waits until there's enough people in the group to buy the item or until the time-limit that the user specifies runs out. We defined and compared two agent behaviours that determined how much the agent offers on behalf of the user:

Goal-oriented – the agent attempts to join a group by offering to pay 10% less than the recommended retail price (RRP). This behaviour can be considered predictable because if one must purchase a product, then one should also be willing to pay up to the retail price. In our study, we emphasised to the participants that the items *must* be bought to ensure they put importance to them[2]. Additionally, an agent that behaves in a goal-oriented manner should also have a higher chance of finding a group that would be completed, as such an agent would make high offers. We also consider that being goal-oriented is a behaviour of how a user expects a typical machine to work.

Fair – the agent attempts to join a group using an offer that can be considered fair (e.g. an equal share of the wholesale price). In contrast to the goal-oriented agent, the fair agent behaves in possibly a less expected way, as it only tries to join a group where the user will pay fairly. Because of this, its chances of joining a group is lower than a goal-oriented agent. However at the same time, it acts more *human-like*, as people put importance to fairness based on the ultimatum game.

In addition to setting up agents as described above, participants were also required to configure the agent's autonomy level for ordering the items, which also determined whether they received notifications or not. We defined three autonomy levels based loosely on previous research (Alan et al., 2016a):

Confirm before buying (low) – when the agent finds a group, it asks the user first whether to buy the item or not. The offer expires after 1-minute;

[2]The interviews confirmed that this was indeed the case.

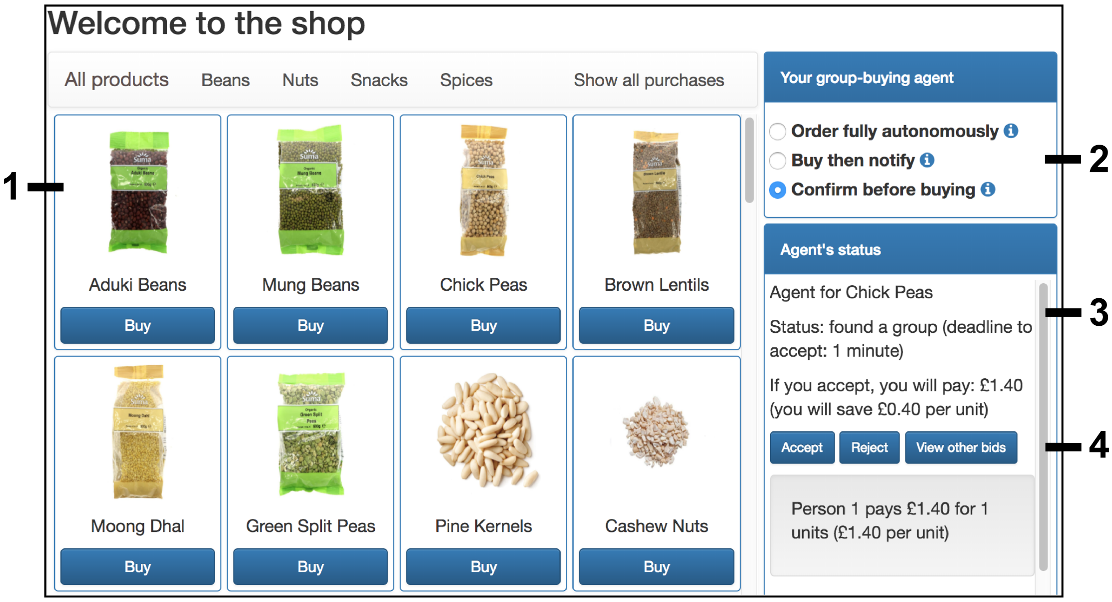

Figure 4.1: A screenshot of the shopping interface, showing the list of products (1), autonomy setting of the agent (2), the agent status, showing the savings and how much the user will pay (3) and the options to either accept or reject the offer and view how much others are paying (4).

Buy then notify (medium) – the agent buys the item when it finds a group and then notifies the user;

Order automatically (high) – the agent buys the item when it finds a group without sending a notification to the user.

An audio alert would trigger when the user receives a notification. Each notification included the price the agent offered and what savings would be made. Participants *could view* how much each person in the group pays in total and per unit upon pressing the "view bids" button. We deliberately left this information "behind a click" so that we could observe how frequently participants would look at it. Participants also received a notification if the agent failed to find a group, or if the time expired for accepting or rejecting an offer.

4.2.2 Study design

We designed our study with the intention of mimicking a real-life situation that involves shopping. In some situations, shopping can be considered a leisure activity (e.g. some people love shopping for clothes, or investing time in finding the best bargain). In contrast, others think of it as a 'chore' that consumes one's limited amount of time (Dholakia, 1999). For instance, doing grocery shopping on the side of full-time work and child care responsibilities has an opportunity cost of not being able to do more pleasurable activities. At the same time, shopping is important and failing to do so also has a consequence e.g. of not being able to eat dinner. In our study, we focus on the type of shopping that can be considered a chore, as this is where we expect autonomous shopping agents to be more appropriate. To simulate this real-life situation in the lab, we took the opportunity of adapting a study method from prior work (Verame et al., 2016). In more detail, participants in our study have the choice to freely switch between two tasks: a shopping task (Section 4.2.4) and an image-tagging task (Section 4.2.3), under a limited period of time. In such a context, if participants perform the shopping task, they lose the chance of earning money through the image-tagging task, with an associated opportunity cost caused by the time-limit. In this case, if a user chooses to invest time interacting with the agent, doing so would cost the user time and effort, much like in the real-life situation we discussed.

We also devised a reward mechanism inspired from behavioural economics that would engage participants in both tasks by linking financial incentives to their performance (Ariely et al., 2008). Participants were allocated a *budget* to complete the shopping task and they get to keep whatever is left of that budget. The budget given to them was equivalent to the sum of all the items being bought at retail price, meaning that if they do not engage in group buying, they would be left with zero. Furthermore, a penalty was incurred for not buying an item, so that if they did not buy anything they would also get zero. This was done to simulate the inconvenience of not managing to buy something that one needs. However, participants could top-up their budget by completing the image-tagging tasks. By so doing, we again aim to simulate a real-life situation, where one can earn money through their day-time job but would also need to spend it for the usual food shopping. Participants were compensated for their time and were also given the remainder of their budget as an additional reward.

While we aim to simulate realism in our study design, our goal is to observe the differences between two alternative conditions in a carefully controlled setting. Specifically, *everything else being equal*, we aim to investigate whether designing an autonomous agent to be fair or not affects people's willingness to adopt it. Most of the previous studies that investigated adoption of autonomous agents have employed an in-the-wild approach to achieve realistic results (e.g. Bourgeois et al. (2014); Mennicken et al. (2016); Yang and Newman (2013)). While we agree with such an approach and believe that it is important to run field trials, it would be difficult, if not impossible, to compare the two agents' designs through a field trial, because of the lack of control over external variables beyond our interests (e.g. when food is required, shopping enjoyment, etc.)

4.2.3 Image tagging

For this task, participants played an image tagging game, where they were asked to guess 3 out of 6 tags describing a particular image (see Figure 4.2). For successfully guessing 3 tags, they would earn 25p, which was added to their budget. There were 60

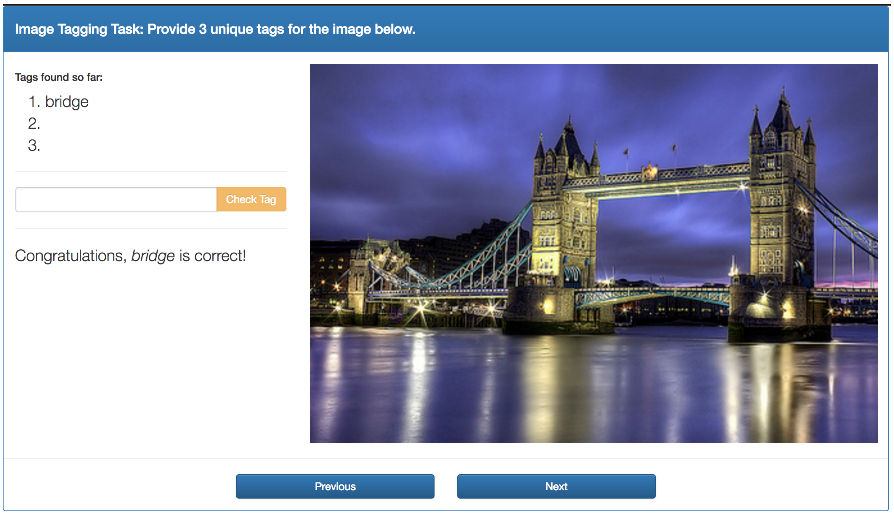

Figure 4.2: A screenshot of the tagging game.

images available for participants to try and they can skip and also re-attempt images if they desire. The pictures and tags were taken from Flickr[3].

4.2.4 Shopping task

For the shopping task, participants were required to buy 8 items from a custom-made website (Figure 4.1). Participants had two options to buy an item (Figure 4.3):

Shop – participants could buy a product as if it was done through a regular shop and pay the RRP;

Wholesale (group buying) – participants can join other people to buy wholesale products and share both the cost and content. From the custom-made website, all wholesale products were sold in packs of 2 units. Items bought this way will always be cheaper than buying from the shop at the RRP, but it may not always be possible to find a group (hence risk of wasting time not being able to buy anything). We developed a "simulation" of how such a market could work, with agents that buy and sell, to mimic people acting in the market (detailed in Section 4.2.5). However, it should be emphasised that our aim was simply to create a situation that would appear plausible from the point of view of the participants and would give them a feeling of impression for a goal-oriented or fair agent, rather than *aiming* for a realistic simulation of a market. It is also worth noting that an important element of our study is ensuring that participants did not know they were playing against a simulation to ensure ecological validity. Indeed, participants were not informed that they are playing in a simulation, and if they asked, they were told that such information cannot be revealed until the end of the experiment.

Participants received an initial budget of £12.50, which is equivalent to the sum of the RRP of all the items that need to be bought. This was done to ensure that participants

[3]https://www.flickr.com/

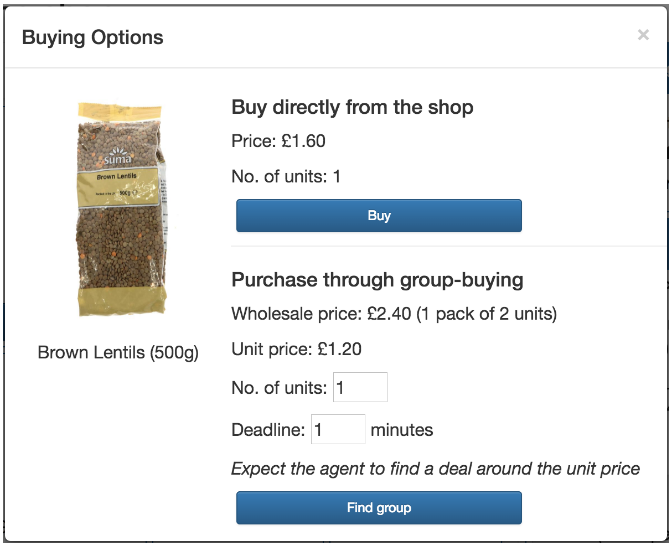

Figure 4.3: A screenshot showing the two options to buy an item. Participants can change how many units they want to buy and how long they want the agent to search (deadline).

have enough money to buy all the items at the most expensive price. If participants do not buy an item from the shopping list, they will receive a penalty. The penalty is equal to twice the amount of the item's RRP. This was done so that if they don't buy any of the items, their reward for completing the task will become zero, which adds emphasis to the importance of completing the shopping task.

4.2.5 Market simulation

The market worked as follows:

1. The script loops through all the available products in the database.
2. For each product, an agent is generated at random. The agent will provide an offer based on its behaviour.
3. The created agent is then added into a queue of agents for that product. If a user creates an agent, it is also added into such a queue.
4. The script loops through the queue of agents. For each agent, it checks whether there is an existing group that it can join. If none, it will start its own group.
5. If the agent does not find a group or if the agent is in a group that does not completed after a number of iterations, the agent leaves the market. It is essentially removed from the simulation.

For a more detailed information of the market simulation, the code is available at: https://bitbucket.org/jhim_verame/groupbuying.

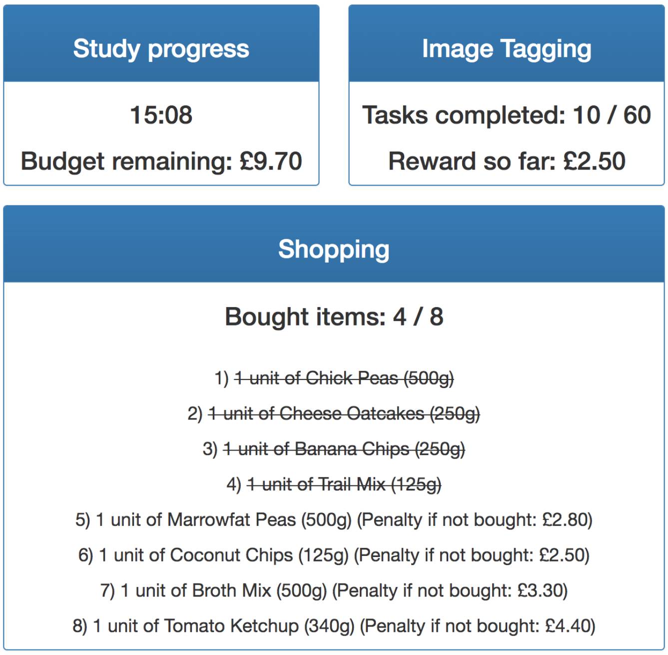

Figure 4.4: A screenshot of the study progress page.

4.2.6 Hypotheses

We are particularly interested in how the agent's behaviour affects participants' willingness to engage in group buying. Firstly, given that an agent, regardless of its behaviour, is capable of helping them save money, participants should be inclined to use it regardless of the uncertainty associated with using it:

- *H1* – Participants will buy more items through group buying than through the shop.

Secondly, having an agent that behaves in a goal-oriented manner in terms of how it finds offers will be less preferred to an agent that takes fairness into account:

- *H2* – More offers will be rejected in the *goal-oriented* condition than in the *fair* condition.

Finally, participants will want to have control over the purchases the agent makes:

- *H3* – More participants will opt for the low autonomy level (*confirm before buying*) than the other two autonomy levels.

4.2.7 Participants

A total of 20 participants (6 males) (with age in years M = 21, S.D. = 8.62) and 10 per condition were recruited through a departmental subject pool. All participants were fluent in English and 17 participants were university students from a variety of disciplines (Psychology, Accounting, Medical Sciences and Engineering). The other participants were of a working professional background (social worker, research assistant and financial planner).

4.2.8 Method

At the beginning of each experiment, participants were assigned in one of the experimental conditions. Conditions were determined prior to running the experiments to avoid mistakes during the instruction phase. Conditions were also set in alternate (e.g. fair then goal-oriented then fair again) and to avoid bias, each day started with a different condition (e.g. on Day 2, the series of experiments started with goal-oriented first). Before each experiment began, participants were asked to read the instructions and sign a consent form. The researcher then showed them the three different parts of the system. The set up involved a laptop plugged into an external 22-inch monitor. On the additional display, a progress page was displayed, showing information about the study such as the time and budget remaining, number of tagging tasks completed and the shopping list (Figure 4.4). The other two pages were displayed on the laptop screen. Participants completed a short trial version to help them gain familiarity with the system before the actual trial to ensure they do not misunderstand how to complete the study. Their experience of the short trial was similar to the actual condition, but with different products in the shopping list and different images on the tagging game. More importantly, they also experienced a similar agent behaviour as they would experience in the actual study. After this short trial, they then completed the actual trial for 20 minutes. This was then followed by a questionnaire about their demographics and then a short interview. The whole experiment lasted around 45 minutes. As detailed in Table 4.1, participants' rewards were linked to their performance in the study, and it ranged from £2.13 to £12.96. In addition, all participants were rewarded a base compensation of £5 for taking part in the study.

4.2.9 Data collection

Data was collected through a combination of quantitative and qualitative techniques. We were particularly interested in the following dependent variables:

- *Items bought through agents* – the percentage of items bought from wholesale through an agent out of all items bought;

Table 4.1: Descriptive statistics showing the means (M) and standard deviation (S.D.) for each participant.

	Fair		**Goal-oriented**		**Overall**	
	M	*S.D.*	*M*	*S.D.*	*M*	*S.D.*
bought through agents (%)	91.25	13.24	100.00	0.00	95.62	10.16
no. of agents created	9.60	3.31	22.80	21.19	16.20	16.24
rejected agent-offers (%)	13.61	17.33	50.90	21.86	32.25	27.10

- *Rejected agent-offers* – the percentage of offers made by the agent that users rejected out of all the offers made by the agents. If an agent's autonomy level is medium or high, its offer counts as accepted. If an agent's autonomy level is low and the participant did not respond to a found offer in time, the offer counts as rejected;

- *No. of agents created* – the number of times that participants attempted to buy from wholesale. This is also the same as the number of agents created.

Moreover, participants were observed by a researcher throughout the study and interviewed at the end, to clarify their actions during the sessions. Each interview lasted approximately ten minutes and was audio-recorded. Interviews were transcribed and analysed by two researchers.

4.3 Results

The study observations and interviews indicated that most participants could relate to the concept of group buying, some having already had real life experience of it, and others feeling encouraged to try it as a consequence of the experiment. Moreover, participants were sensitive to the financial incentives, presented through the budget mechanism; they were aware of the limited duration of the experiment, and of the need to balance the shopping and tagging tasks: they all mentioned trying to earn money by saving as much as possible on the shopping, and by tagging as many images as possible. Most participants also mentioned not feeling pressured in the study, enabling them to explore their options especially for doing the shopping task. Some participants contemplated the possibility that there may be people who are doing the experiment at the same time, suggesting that it was not obvious they were playing in a purely simulated market. Table 4.1 provides a summary of the quantitative data.

4.3.1 Usage of the agents

All of our participants attempted to purchase items through group buying by creating agents. In particular, our participants created a total of 324 agents and 154 (95%) of the items were bought through agents. Independent-samples t-tests of the agent behaviour did not reveal any significant differences on the number of created agents and percentage of items bought through agents.

During the interviews, our participants mentioned that they went with the group buying option since it was cheaper than buying from the retail store. Only four participants purchased from the retail store (for a total of 6 products). They explained that they bought from retail store when they felt that they did not have time to set up new order or when the agent couldn't find a deal.

4.3.2 Setting the autonomy levels

A total of 311 (96%) agents were set to low autonomy whereas the remaining 13 (4%) were set to medium autonomy. No agents were set to the highest autonomy setting. In terms of the number of participants, all but two chose the low autonomy option for their agents. They chose this option as they wanted to make sure that they were happy with the price the agent found before making the purchase. They *"wanted to have the control"* to accept good offers or reject those that they may not like. The two participants (both from the *fair* condition) who chose 'buy and notify' (medium autonomy) shared similar reasons for choosing this option. They did not want to waste time checking the agents' offers since they knew that the price the agents would find would always be cheaper than the retail price. They chose to focus on earning money by doing more of the image tagging task instead.

4.3.3 Accepting and rejecting offers

As shown in Table 4.1, 32% of all agents that found a deal were rejected. In addition, an independent-samples t-test revealed a significant effect of agent behaviour ($t(18) = 4.23, p < 0.05$) on the percentage of agents rejected. There is a higher percentage of rejected agents in the *goal-oriented* condition ($M = 50.90, S.D. = 21.86$) than in the *fair* condition ($M = 13.61, S.D. = 17.33$). Figure 4.5 shows the means comparison of the percentage of rejected offers, with the 95% confidence intervals, for this analysis.

Our analysis of the interviews also revealed our participants' reasons for rejecting and accepting offers. Most participants responded to offers based on how much differently they are paying in comparison with the other person in the group. They would often accept offers if they were paying the same amount as or less than the other person.

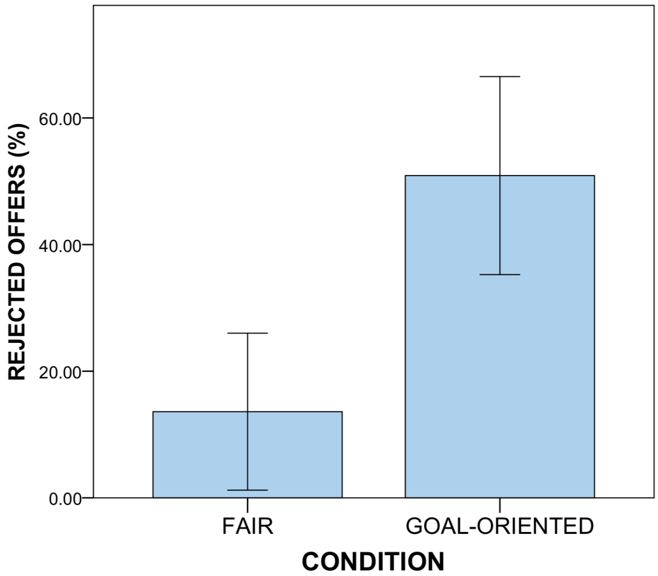

Figure 4.5: Means comparison for rejected offers between the *fair* and *goal-oriented* conditions. The error bars represent the 95% confidence intervals.

Otherwise, if there was a large discrepancy, they would often reject and try to set up agents again. In contrast, those who did not look at what others were paying only took into account whether the price differed from the retail price. A few participants fully delegated the purchasing decision to their agents. One participant also felt that that the risks were too high to reject offers: *"If I reject, I might not able to find a group."*.

4.3.4 Weighing up on fairness

An important variable in determining whether participants paid attention to fairness is if they checked how much the other person in the group is paying. From our usage logs, there were 14 participants (all 10 participants from the *goal-oriented* condition and 4 participants in the *fair* condition) who viewed what the other person in the group was paying. Moreover, a Welch's t-test also revealed a significant effect of bids being viewed ($t(16.65) = -5.25, p < 0.01$) on the percentage of agents rejected. There is a higher percentage of rejected agents when the bids were viewed ($M = 43.47, S.D. = 24.58$) than when they are not ($M = 6.09, S.D. = 6.77$). Figure 4.6 shows the means comparison of the percentage of rejected offers, with the 95% confidence intervals, for this analysis.

From our interview data, 13 participants explicitly mentioned caring about fairness. In the *fair* condition, only 4 participants mentioned caring about fairness. For example, one participant mentioned that because the same product was being shared, they shouldn't have to pay more than the other person. Another participant expressed that he *"didn't want to get into an unfair situation"*. In fact, one participant told us that he rejected offers where he was paying more than the other. The majority of the other six participants who did not care about fairness in this condition explained that as long

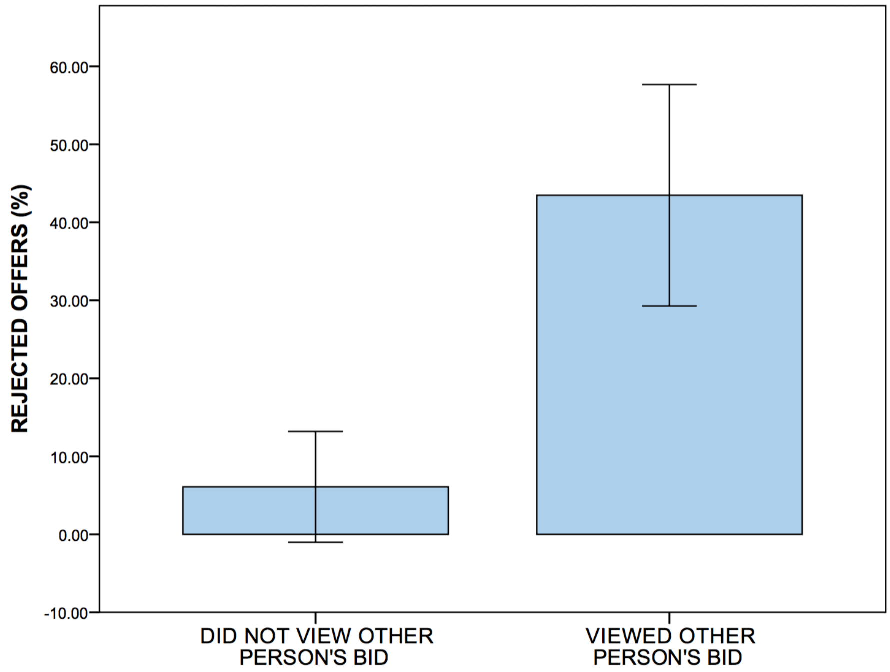

Figure 4.6: Means comparison for rejected offers between viewing and not viewing the other person's bid. The error bars represent the 95% confidence intervals.

as they were paying less than a retail store and thus saving money, they did not care much about what other people were paying. One participant specifically mentioned that knowing what the other person was paying would not help him to save more money.

Contrary to the above, all but 1 participant in the *goal-oriented* condition expressed concerns regarding the fairness of the offers they received. Among the nine, some felt resentment. For example, one participant commented that *"I was thinking about how i'm paying more. I felt it was unfair, a bit questioning why it is the case"*. Another participant commented that the offers she received saddened her, causing her to ignore the prices and just accept blindly. The one participant who did not care about fairness in the *goal-oriented* condition shared a similar reasoning to those who did not care in the *fair* condition: as long as there was a cheaper price than a retail price, they did not care about other person's price.

Indeed, there were instances that even when the participants felt that the deal was unfair, they would still accept the offer. Participants reasoned that this was because it was still cheaper than the retail option and there was no way for them to change the situation. Some tried to find a better deal by rejecting the offer, however they discovered that the offer stayed the same. They reckoned that by accepting the offer, though it was unfair to them, they would still manage to save some money.

4.3.5 Orientation towards the agents

There is a distinction between the *goal-oriented* and *fair* conditions about our participants' orientation towards the agents they used. All but 1 of the participants in the *fair* condition expressed that they were happy with their agents. They found that the agents were helpful and some even mentioned that their agents were trustworthy. For example, when asked about why he was happy with the agents, one participant replied *"because he did a hard work for me ... I don't know if he choose the first one he sees or the cheaper, he was at least trying to save for me"*. The remaining one participant in the *fair* condition who did not express any positive opinion about their agents reasoned out that he would only be able to trust the agents if the products they buy are of good (physical) quality.

In contrast, most participants in the *goal-oriented* condition (7) have either reserved opinions about their agents or were not happy at all. They felt that they could not trust their agents, especially because they perceived that the agents were not doing their best to find good deals as they were always paying more than the other person. For example, when asked whether she was satisfied with the agents she used, one participant responded: *"Not really .. say if I would need to invest in it or anything, I won't really consider because it is just saving me, like, may be 8p"*. The other 3 participants in the *goal-oriented* condition instead reported that they were happy with the agent because it still saved them some money, albeit the savings being small.

4.4 Discussion

The study presented in this work allowed us to compare two alternative design choices for an autonomous group buying agent: one based on fairness and one based on its orientation to completing a goal. Such a comparison was made possible by the scenario presented in our study, in which participants had to balance the chore of grocery shopping with an alternative task, image tagging, for which they were rewarded financially, mimicking a job. The two tasks were linked through the *budget*: the amount of money in it reflected the combined performance on the two tasks, and defined the financial reward participants received at the end of the study. Moreover, the participants' reports of the strategies they adopted during the studies, as well as their drive towards the financial reward suggest that our experimental design was successful in exposing them to a complex, yet controlled scenario through which we could observe their inclination to delegate the shopping tasks to autonomous agents.

4.4.1 Utilising the agents

All participants reported that they attempted to buy every item through the autonomous shopping agents. In particular, over 95% of the items bought were purchased through the agents. These results confirm H1, in that, participants were engaged in using the group buying agents despite the uncertainties involved. Our interviews further revealed the main factors for such a high uptake of the agents. Firstly, an agent's utility, defined by the possibility of making savings, enticed our participants into using them. Indeed, even in the *goal-oriented* condition where the savings were relatively low (10%), our participants still invested time into setting up agents. This finding aligns with prior work, suggesting that when an agent has a higher utility than a manual option (which in this case is buying from the shop), people adopt it (Smith et al., 1997; Verame et al., 2016). Secondly, another reason reported by our participants for using the agents is that they were doing work on their behalf, aligning with findings from previous studies (Costanza et al., 2014; Alan et al., 2016a). Some of our participants explained that having the agents in the background allowed them to spend time on the image tagging task and earn from that.

4.4.2 Importance of fairness

The statistical analysis of our results revealed that the agent behaviour had an effect on whether an offer is accepted or not. In more detail, a higher percentage of offers were rejected in the *goal-oriented* condition than in the *fair* condition. Furthermore, 70% of our participants viewed the information about what the other person would be paying. Our results also revealed that checking this information had an effect on whether an offer is accepted or not. In particular, a higher percentage of offers were rejected when the bids were viewed than when they were not. In addition to these quantitative results, the majority of our participants explicitly talked about fairness as an important factor of their actions in the study during the interviews. Indeed, both quantitative and qualitative data from the study suggest that having an agent that behaves fairly can affect participants' willingness to adopt them, thus accepting H2. Even when participants in the *goal-oriented* condition chose to accept the agent-offers, the interview data show that most of them were not happy or did not feel they could trust their agents, which gave one of our participants a reason to *"not invest"* in such an agent. Whether this means investing by purchasing or by spending time interacting with the agent, both have negative implications for the adoption of this type of agents.

It is worth noting that though most of our participants seemed to care about fairness, they did not care as much as people who played the ultimatum game. In the ultimatum game, participants give much importance to fairness that they were willing to not receive anything rather than receive unfair offers. In contrast, our participants, regardless of the condition, typically ended up accepting unfair offers, with the justification that the offers were still better than paying the higher retail price. It is possible that such differences may have been a result of the framing of the study, as different framings have been shown to affect people's decision-making (Kahneman and Tversky, 1984). In our study, participants were explicitly told to try and earn as much reward as possible and also to ensure they buy all the items. In contrast, although desirable, people who play the ultimatum game are not explicitly told to earn as much reward as they can (Guth et al., 1982). The differences may have also been a result of how the agents used in our study were designed. In our study, the agents were designed to be simplistic and to not have any learning capabilities, which was realised by some of our participants. This means that even if the offers they make were rejected, the agents do not take into account those responses in future offers. Participants, especially those in the *goal-oriented* condition, who noticed this eventually accepted unfair offers to avoid paying the costlier retail price. In contrast, people who played the version of ultimatum games with multiple rounds often learn to adjust their offers based on the response of the other party (Gneezy et al., 2003).

4.4.3 Retaining control

While there is a high uptake of the agents, the majority of the agents created (96%) were configured with the lowest autonomy setting, which allowed users to confirm before the agent purchases an item. Furthermore, our qualitative data reveal that the most common reason for this choice of autonomy level configuration is for users to have control over the decision of their agents. This confirms H3 and suggests the importance of ensuring users have control over their agent's autonomy, aligning with previous studies (Alan et al., 2014, 2016a).

4.5 Implications

We conclude our work by reporting a number of implications for the design of future autonomous group buying agents.

Feasibility of group buying agents. Our findings highlight that participants found the autonomous agents helpful in doing group buying. Many of them bought items using the agents and some of them even delegated the purchasing decision to the agents. Such findings align with a recent marketing survey, suggesting that shoppers are ready for automated purchases[4]. Based on this, researchers should explore the design of such shopping agents, particularly those that can facilitate group buying for grocery shopping. Indeed, browsing online markets, finding other people, and joining groups require a great

[4] http://www.salmon.com/en/programmatic-commerce-report/

amount of time and effort for users, which could be spent for their daily jobs and other tasks with higher priorities. Having an agent that does all these on behalf of users can greatly benefit them, as our study indicates.

Applying fairness into agents. Our study results indicate that a shopping agent's inclination towards fairness – i.e. whether it is fair or goal-oriented – seems to affect whether people adopt it or not. In particular, having an agent that acts fairly can result in higher uptake than one that focuses on its goal. Based on this, the design of such autonomous agents should therefore have an orientation to fairness, following also on the principle that technology should be designed according to human values (Friedman and Kahn, 1992; Friedman et al., 2008). Indeed, recent work (de Melo et al., 2016, 2017a,b) has already suggested that agents that act fairly may help to enforce fairness in society. Extending this work, our findings suggest that although people are likely to reject an "unfair" (or goal-oriented) agent, they may still use it if it still provides some utility.

Furthermore, while current research on mechanism design aims to guarantee certain fairness properties (e.g. Roos and Rothe (2010); Cohler et al. (2011); Chan and Chen (2016); Alijani et al. (2017)), our results indicate that fairness does not have to be perfect for people to adopt the technology. The assumptions on the rationality of agents or their owners can be relaxed to enable the creation of mechanisms that do not make too many assumptions about the form or distribution of agent preferences (or types).

Making agent performance clear. Our results suggest that even though the agents employed in this study did not seem ideal to our participants, especially in the case of the goal-oriented agents, in which our participants were paying more than the other person and the price offered by the agent was not significantly lower than the retail price, participants still accepted their offers. The main factor for this is that our participants had a sense of what the agents could do (e.g. finding a price cheaper than RRP). Prior work has suggested that increasing transparency of how autonomous systems work can lead to improved user adoption (Rodden et al., 2013; Yang et al., 2015b). However, increasing transparency may result in higher workload for the users (to process the information) and even the risk of causing frustration, if users find it difficult to understand the underlying technology. Instead of displaying how such a system works, designers can focus on the agent outcome, setting users' expectations for the agent's level of service. Such an approach may encourage users to adopt the autonomous agents, without requiring them to fully understand the agent system's technical details.

Enabling flexible autonomy. Our results also suggest that flexible autonomy can be beneficial for users to help them interact with autonomous group buying agents, aligning with prior work (Alan et al., 2014, 2016a). As we have seen in our study, whereas the majority of the participants chose the lowest autonomy setting, some participants still created their agents with the medium autonomy level. The majority of the participants preferred having control over their agent making decisions for purchase. The others who chose the medium level preferred saving time and earning money by doing other tasks. A flexible autonomy level can be beneficial to users by allowing them to choose different levels according to their preferences and priorities.

4.6 Limitations

Since our study is based on a controlled lab setting, the parameters in the simulation used in the study were defined to accommodate that, which might not fully reflect how group buying agents work in real life. For example, the agents in the study found deals in a relatively short amount of time (e.g. within minutes), which is something that may not always happen in real life (it could take hours or even longer). Future research should investigate the feasibility of shopping agents beyond lab studies and into the wild. In addition, we opted to not inform participants that the study was purely simulated, in an effort to elicit actions from participants that are realistic, similar to what has been done in other studies. While some participants may have thought that the study was not simulated, others may not think that way, which may have potentially polluted the data. However, we note that during the interviews, participants did not talk about the simulated nature of the study as a reason for their actions. In fact, once they were debriefed, none of our participants felt that their behaviour would have changed had they known. Despite this, future studies should look at group buying scenarios with real people, to ensure that participants will behave more naturally. Another limitation is that in real life, shopping decisions are made on a multitude of factors. Although we designed our study in order to reduce these factors as much as possible to enable us to focus on fairness, future studies should investigate whether other factors related to people's shopping decisions can affect the importance of fairness.

Regarding the agents employed in the study, our work is based on two agents that both provide utility (finding a deal cheaper than RRP), where one always find fair deals and the other finds unfair deals. We purposely chose these two agent designs, in order to draw out the importance of fairness by comparing participants' attitudes towards both agents. However, it is important to point out that this may not always be the case. For example, for goal-oriented agents, they can potentially find deals that are also fair for the user. Future work should therefore explore different agent designs.

4.7 Summary

In this chapter, we reported on a lab study designed to explore the importance of personal values in the design of autonomous shopping agents, particularly around group buying for grocery shopping. The study employed a combination of quantitative and qualitative methods, and crucially financial incentives related to the participants' performance in a multi-task setting to achieve higher ecological validity. In more detail, 20 participants were exposed to two alternative group buying autonomous shopping agents: one designed around fairness, and the other around being goal-oriented. Our findings revealed that all participants welcomed the support provided by the autonomous agents, particularly in their ability to guarantee positive outcomes. However, participants were more accepting of agents that take fairness into account. An implication of our work, then, is that the concept of fairness from BGT seems to also transfer to agent-design in the scenario we explored. This suggests that taking into account personal human values help people deal with the uncertainties involved with using autonomous agents.

Indeed, it has been suggested by existing research that incorporating human attributes in agent-design can be beneficial to help users interact with agents. Extending prior work, our study provides empirical evidence that agents that take into account human values are also more likely to be adopted than those that only focuses on achieving a goal. This then addresses R2 of our research questions, in that an autonomous agent should therefore have the characteristics of a human to increase its chances of adoption, especially in situations where there are uncertainties in the performance of the agent. However, we acknowledge that it may not always be possible to incorporate human values in autonomous agents. For instance, it may difficult to embed human values in agents that are purposely designed to help people perform tasks quicker and easier, such as an optical character recognition (OCR) or a grammar correction software. Therefore, an alternative feedback mechanism that can be integrated in the design of autonomous agents would be beneficial. One way to achieve this, as outlined by one of our design implications, is by ensuring that the performance or the level of service that an agent can provide is clear to the users. Based on this, we are interested to see whether displaying the potential or estimated performance of an agent can help users delegate agency to it. In the following chapter, we detail a study that explores this concept.

In addition, we also described a study design, which has been inspired by behavioural economics literature, that enabled us to observe the differences between different agent-designs in a controlled setting. By so doing, it allowed us to investigate which agent-design can encourage people to adopt autonomous (group buying) agents. This then addresses the research question R4, showing that it is possible to design a study that can help us properly evaluate how different agent-designs affects people's inclination to make use of an agent.

Chapter 5

Displaying Confidence Information of Autonomous Agents

The group buying study has revealed that a software agent with human values (e.g. an agent that takes into account its owner's sustainable aspirations or one that considers whether a deal that it finds is fair to its owner) can encourage people to adopt it and also delegate agency to it. However, as mentioned in Section 4.7, it may not always be possible to incorporate human values in agent-design. Without the ability of designing human values in agent design it is important to investigate alternative interaction mechanisms or interface design that can help users in dealing with the uncertainty of autonomous agents. Specifically, in this chapter, we further explore whether displaying the agent's level of service can be useful in encouraging users to utilise autonomous agents.

This chapter covers R3 of the research questions, as stated in Section 1.3. Furthermore, the study method described here is similar to the group buying study, which is aimed at addressing R4. We start the chapter by describing the motivation of this work. This is then followed by detailing the study design. We then outline the findings and provide discussion points. Finally, we provide design implications for future autonomous agents and summarise the whole chapter.

5.1 Motivation

One of the design implications from the group buying study is that an agent's performance or its level of service should be transparent to the users. Having this information allows users to set an appropriate expectation. For example, in the group buying study, users were aware that although they cannot guarantee that the agent will find a group, they know for a fact that when it does, it will always find a deal cheaper than the RRP. Knowing this information allowed participants to trust the agent, even to the extent of letting the agent buy automatically on their behalf. However, as can be observed in our studies, it is not always possible to provide an accurate information about the agent's level of service. In some instances, it may only be possible to provide an estimated probability of what an agent can achieve. Based on this, it is important to understand whether displaying the estimated probability of an agent's level of service, or its *confidence information*, can also help users in delegating agency to autonomous agents.

Indeed, as discussed in Section 2.1.4, prior work has suggested that displaying confidence information can increase a user's awareness of the ability of autonomous agents (Beller et al., 2013; Helldin et al., 2013). However, these studies lack ecological validity and assumed users will always interact with the autonomous agent. Moreover, they are often focussed on user performance rather than adoption and agency delegation. Against this background, in this chapter we report a lab study designed to investigate whether people can make sense of confidence information and whether displaying such information can improve user adoption of autonomous agents. Our results demonstrate that confidence information encourages the usage of the autonomous agent we tested, compared to situations where such information is not available. We describe the details of our study in the next section.

5.2 User study

A user study was designed and conducted to test the effectiveness of the confidence information of an autonomous agent. Specifically, we wanted to examine whether the confidence information affects a user's decision to interact with an autonomous agent helping users in an activity that can be considered mundane or common to various people (e.g. students, office workers, researchers). So we looked for an example autonomous agent around which we could set up a credible scenario to play out in a lab study, and which could be related to tasks which would be natural for a population of university students and administrative staff. We chose tasks and agents related to common *textual document* activities, such as typing up handwritten notes and proofreading text. We chose an activity outside the domain of food to further reduce the number of factors that may influence people's actions in the study. Furthermore, we chose these tasks as we had access to the ground truth. This allowed us to automatically check the correctness of each submission and provide immediate feedback to participants about their performance. The study has been approved by the Southampton Ethics Committee (ref: 19339).

In the following subsections we first describe the autonomous agent and then detail the tasks used in the study.

5.2.1 OCR agent

We designed an agent which automatically recognised handwritten text and converts it into typed text – an example of an Optical Character Recognition (OCR) software. OCR applications are widespread and likely to be familiar (at least conceptually) to most of our participants.

The agent processes one document at the time, taking roughly 30 seconds. Participants were told that the agent may make mistakes, which would need correcting. These mistakes were incorrect type outs of characters that may look similar to other characters (e.g. the letters *n* and *h*). The OCR processing goes on in the background, autonomously. Upon the agent completing the task, a sound goes off indicating the availability of the results to the user (similar e.g., to receiving an incoming email message). Furthermore, as the agent completes tasks, they get added to a queue regardless of whether the user attended or ignored the previously completed task(s), somewhat similar to an email inbox.

Because of the possible mistakes made by the agent, users are required to 'review' or 'accept' the completed agent tasks. More specifically, for any completed agent task available in the queue, users had three options about how to deal with it:

1. *Review* – participants can view the task result to check and correct any errors.
2. *Blindly accept* – participants can blindly accept the agent result without reviewing it, essentially *fully accepting* the agent automation.
3. *Ignore* – participants can also opt not to review the completed agent task and just leave it in the queue. These tasks can be reconsidered at any later moment.

To ensure that the performance of the agent was consistent, as this could affect how participants use it, we adopted a Wizard-of-Oz approach. The agent was actually artificial, in that all the handwritten documents had been originally typed in by a researcher and errors were introduced in a controlled manner, to simulate 5 different levels of confidence on various documents. In particular, the five levels of confidence are: *very low, low, medium, high* and *very high.* Each level is related to the number of errors in the task (see Table 5.1 for error rates). Participants were not informed that there is no real OCR agent, and if they asked, they were told that such information cannot be revealed until the end of the experiment.

Table 5.1: No. of documents with at least one error per confidence level. For example, for agent tasks with *medium* confidence, half of them had at least one mistake. Note that because of randomness the *high* confidence level has fewer errors than *very high*.

Confidence levels	Very high	High	Medium	Low	Very low
No. of documents with at least one error	1/6	0/6	3/6	5/6	5/6

5.2.2 Alternative tasks

To implement a choice situation where interacting with the agent would create an opportunity cost, we defined an additional task in addition to the **agent task** described above. This task, which we call the **manual task**, involves correcting the grammar of a 6-line long paragraph typed in English, checking for singular or plural agreement (*is, are, has, have*) and also for commonly mistaken possessive terms (*they're, their, it's, its*).

Although the agent helped users complete the transcription task, we wanted to make the work of reviewing agent tasks require more effort than completing the manual task. This was because in real-world situations, monitoring work completed by autonomous agents would require users to invest time that could otherwise be spent on doing other activities, especially at the beginning, when they have little experience with the agent. For this reason, the documents processed in the agent task were written in a foreign language which would not be familiar to our study participants: Filipino[1]. This is to ensure that if users were to review the task, they would actually be comparing the handwritten and typed text. Such may not be the case if the manuscripts were in a common language (e.g. English or Spanish), where users may simply check the spelling of the digitised text. In the next section, we detail the interactive system used in our lab study.

5.2.3 User interface

We designed and developed an interactive system which simulated the scenario explained in the previous sections. Figure 5.1 shows the interface, which is divided into four main panes:

Dashboard. The dashboard contains statistical information about a user's status during the study. It displays the number of correct and submitted manual and agent tasks. Furthermore, the dashboard also shows the current reward and time limit.

Manual task switch. Allows users to switch to the manual task.

[1] Anyone familiar with this language was excluded from our sample.

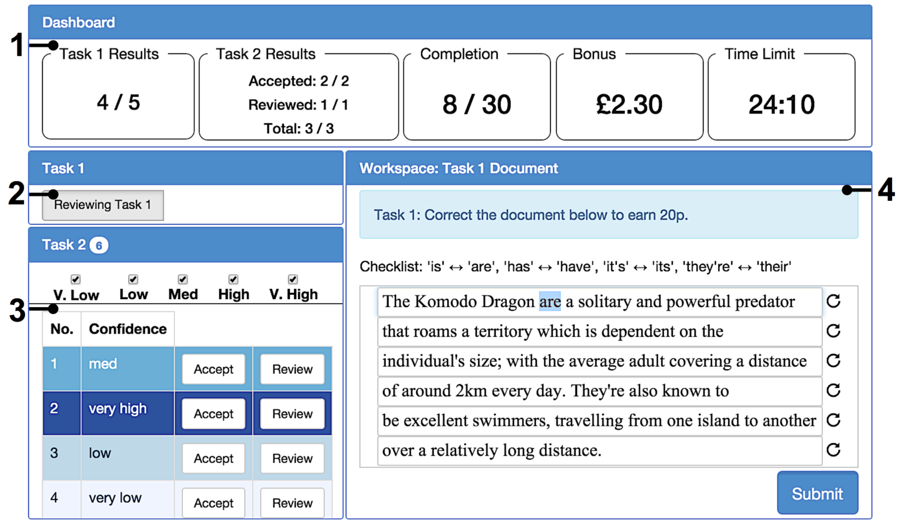

Figure 5.1: Screenshot of the interface, showing the dashboard (1), manual task switch (2), notification panel (3) and workspace (4). In the workspace, the highlighted word *are* must be replaced to *is*.

Notification panel. This panel shows agent tasks as they become available, where each row corresponded to one agent task. The *Review* button allows users to view the agent task, which will be shown in the current workspace. The *Accept* button allows users to blindly accept agent tasks. Each row also contains information about the confidence of the agent for that task. The rows are coloured according to the associated confidence (the higher the intensity, the higher the confidence). Furthermore, a filter function is available to help users filter the tasks based on the different confidence levels.

Workspace. The workspace shows the current task being worked on. For example Figure 5.1 is showing the manual task, whereas Figure 5.2 shows an agent task with very low confidence.

5.2.4 Design

A 2×3 between-subjects study design was employed[2]. The *confidence information* was manipulated as an independent variable (IV), through the following conditions:

- Confidence – participants were able to see the agent's perceived confidence for each of its completed task.
- No-confidence – the confidence information was omitted. In addition, the agent tasks in the notification panel were not coloured.

[2]A within-subject design was not possible because of the learning effect associated with the confidence information and also the types of errors in both tasks.

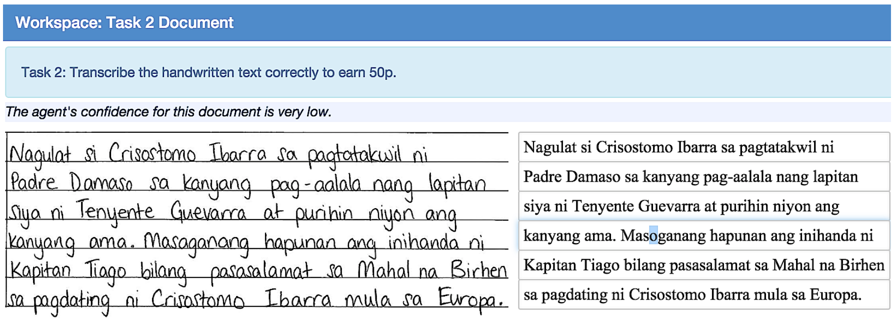

Figure 5.2: An example of an agent task. In this example, the highlighted letter *o* (4th line on the right side) must be replaced by the letter *a*.

We also manipulated the *incentive scheme* as an IV to validate the methodology we used in the study, with 3 conditions:

- No-incentive – participants were paid £6 for their participation, regardless of their performance in the study.
- Agent-incentive – participants were paid 50p for submitting an agent task without any mistakes and 20p for correcting all the grammatical mistakes in a manual task.
- Manual-incentive – participants were paid 20p for submitting an agent task without any mistakes and 50p for correcting all the grammatical mistakes in a manual task.

In the *agent-incentive* condition, the choice of payment reflects the amount of effort and time required to complete each of the tasks. During the pilot runs, manual tasks were completed around 20 seconds in average, whereas the agent task took around 50 seconds. In short, the agent task took $50/20 = 2.5$ times more time (and therefore effort) as the manual task. For the *manual-incentive* condition, we reversed the incentives used in the *agent-incentive* condition. This was done to double-check whether the amount of incentives used in the *agent-incentive* were sufficient to motivate participants to choose one task more than the other. Furthermore, the *manual-incentive* condition should also negate or reduce the impact of factors other than the monetary reward that would affect users choosing the agent task. The next section details our hypotheses.

5.2.5 Hypotheses

We are particularly interested in how the confidence information affects participants' inclination to review or blindly accept agent tasks. The confidence information should make it possible for participants to know which agent tasks require lower effort (the ones with higher confidence). So we hypothesised that:

H1a – When confidence information is displayed, participants will use the agent **more**. In particular, they will complete (i.e. review or accept) a higher number of agent tasks than when the confidence information is omitted.

H1b – Participants will complete more agent tasks with high confidence than agent tasks with lower confidence levels.

Secondly, confidence information should also inform users when the agent can be relied upon and when users need to intervene:

H2a – When confidence information is displayed, participants will rely more on the agent – i.e. they will blindly accept more agent tasks than when the confidence information is omitted.

H2b – Participants will accept more agent tasks with high confidence than agent tasks with lower confidence levels.

Additionally, we expect that our experimental methodology would affect the decision of users in choosing between completing the manual and the agent task. In particular, the different financial incentives imposed should influence users about which of the two tasks they should complete more. If so, this would validate our methodology. Our final hypothesis therefore is:

H3 – Participants in the *agent-incentive* condition will complete more agent tasks than manual tasks. Additionally, participants in the *no-incentive* and *manual-incentive* conditions will complete more manual tasks than agent tasks.

5.2.6 Participants

A total of 60 participants (39 female, 21 male) took part in the study, 10 per condition and 59 of these were members of the university: PhD, Masters students and undergraduate students, from a variety of disciplines (including Engineering, Languages, Business and Management, Law, Health and Social Sciences, and Geography). One participant works for the local council in data management for schools. The ages of these participants ranged from 18 to 43 years old ($M = 23.20, SD = 5.43$). As discussed above, the participants we recruited are educated to above average levels, but the tasks defined in our study are suitable for them.

5.2.7 Method

At the beginning of each experiment, participants were assigned to one of the experimental conditions. Conditions were determined prior to running the experiments to avoid mistakes during the instruction phase. Conditions were also set in alternate (e.g. *confidence* then *no-confidence*) and to avoid bias, each day started with a different condition (e.g. on Day 2, the series of experiments started with *no-confidence* first). Before each experiment began, participants were provided instructions about the study. In particular, they were asked to complete up to 30 tasks in total within 30 minutes as accurately as they could. Crucially, participants were given the freedom to select whichever type of task they want to complete and were free to switch from one task to another at any given point in time. Indeed their selection of tasks was a key measure to quantify their inclination to use the autonomous agent. Participants paid based on performance (*agent-incentive* and *manual-incentive* conditions) were told that there is a limit of £10 to earn. The study will end if they reach this amount, run out of time or if they complete 30 tasks. Furthermore, participants in the *confidence* conditions were told that agent tasks have associated confidence levels. Details about how the confidence information was formed were not revealed to the participants. After these instructions, participants completed a 5-minute training period to help them gain familiarity with the system before starting the actual trial. Participants were shown how to switch between the two tasks during this training period. This is to ensure that they would not misunderstand how to complete the study, such as thinking that they need to complete all Task 1 documents first before doing Task 2 documents[3].

5.2.8 Data collection

Data was collected through a combination of quantitative and qualitative techniques. The quantitative analysis was used to check for any statistical significances. Instead, the qualitative analysis was used to gain in-depth understanding of participants' experience during the study. The system automatically measured the following dependent variables:

Agent tasks completed – the proportion of agent tasks completed out of all completed tasks (the combination of reviewed and blindly accepted);

Agent tasks blindly accepted – the proportion of completed agent tasks that were not reviewed by the users out of all completed tasks;

Time – the average time taken (in seconds) for participants to complete the tasks, which can be interpreted as the amount of effort spent by participants;

[3] The interviews confirmed that participants understood that it was possible to switch between the two tasks.

Reward – the final reward received (in £) for participants in the *agent-incentive* and *manual-incentive* conditions;

Correct submissions – the proportion of tasks completed correctly out of all completed tasks;

Completed agent tasks per confidence level – the proportion of completed agent tasks by the users for each confidence level out of all completed agent tasks;

Blindly accepted agent tasks per confidence level – the proportion of blindly accepted agent tasks for each confidence level out of all blindly accepted agent tasks.

Moreover, participants were observed by a researcher throughout the study and interviewed at the end, to clarify their actions during the sessions. Each interview lasted approximately five minutes and was audio-recorded. Interviews were later coded through open codes for each experimental condition, then grouped in categories altogether through thematic analysis (Braun and Clarke, 2006). Open coding was completed per condition to identify main themes within each condition. Then, axial coding was completed for open codes across all conditions, to find the main themes for the whole study.

5.3 Results

The following subsection reports the results of the analysis of our quantitative data.

5.3.1 Quantitative analysis

A total of 1088 manual tasks were completed and 768 of those were correct (70.59%). For agent tasks, 621 were completed with 503 correct (81.00%). Furthermore, 126 of the completed agent tasks were blindly accepted (20.29%) and 99 of those blindly accepted tasks were correct (78.57%). In more detail, there were 50 completed agent tasks in the *no-incentive, no-confidence* condition, 94 in the *no-incentive, confidence*, 184 in the *agent-incentive, no-confidence*, 194 in the *agent-incentive, confidence*, 28 in the *manual-incentive, no-confidence* and 66 in the *manual-incentive, confidence*.

Proportion of agent tasks completed. A two-way ANOVA revealed a significant effect of both confidence information ($p < 0.05$) and incentive scheme ($p < 0.001$) on the proportion of agent tasks completed by participants. No interaction effects were statistically significant. When confidence information was displayed, participants completed a higher proportion of agent tasks. A post-hoc Tukey test on the incentive schemes revealed that a higher proportion of agent tasks were completed in the *agent-incentive*

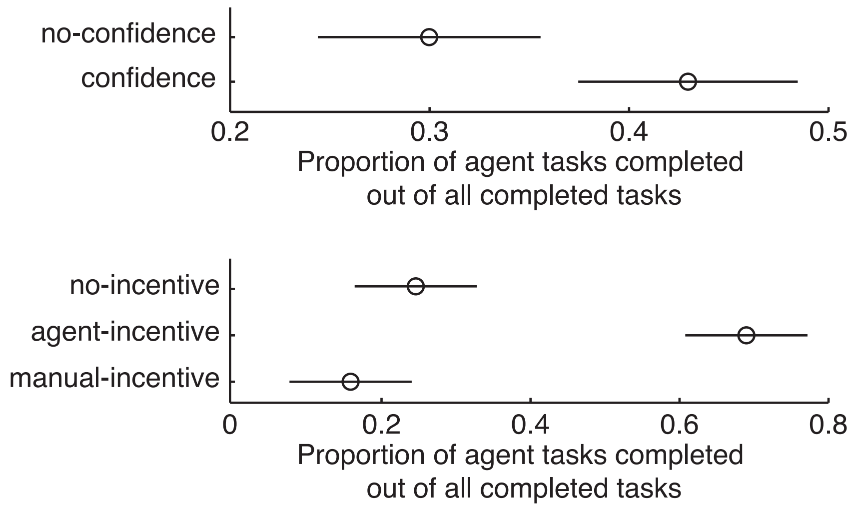

Figure 5.3: Means comparison for agent tasks completed across different displays of confidence information (top) and incentive schemes (bottom), with the 95% confidence bars (Tukey-HSD).

condition ($M = 0.69, SD = 0.24$) than the *no-incentive* ($M = 0.25, SD = 0.25$) and *manual-incentive* ($M = 0.16, SD = 0.18$) conditions. Figure 5.3 shows the means comparison of the proportion of agent tasks completed, with 95% confidence intervals (Tukey-HSD), for this analysis.

Proportion of agent tasks blindly accepted. A two-way ANOVA revealed a significant effect of confidence information ($p < 0.05$) on the proportion of agent tasks blindly accepted by participants, with a higher proportion of tasks being blindly accepted in the *confidence* condition. No statistically significant differences were found based on incentive schemes, nor interaction effects.

Proportion of agent tasks completed per confidence level. A one-way ANOVA revealed a significant effect of confidence level ($p < 0.001$). A post-hoc Tukey test revealed that there were significantly more completed agent tasks with *very high* confidence level ($M = 0.31, SD = 0.21$) than agent tasks with *medium* ($M = 0.19, SD = 0.07$), *low* ($M = 0.17, SD = 0.10$) and *very low* ($M = 0.19, SD = 0.07$) confidence level. Figure 5.4 shows the means comparison of the proportion of agent tasks completed per confidence level for all confidence levels.

Proportion of blindly accepted agent tasks per confidence level. A one-way ANOVA revealed a significant effect of confidence level ($p < 0.001$). A post-hoc Tukey test revealed that there were significantly more blindly accepted agent tasks with *very high* confidence level ($M = 0.55, SD = 0.27$) than agent tasks with *high* ($M = 0.30, SD = 0.18$), *medium* ($M = 0.05, SD = 0.09$), *low* ($M = 0.03, SD = 0.05$) and *very low* ($M = 0.07, SD = 0.26$) confidence level. The *high* confidence agent tasks were also blindly accepted significantly more than agent tasks with *medium*, *low* and

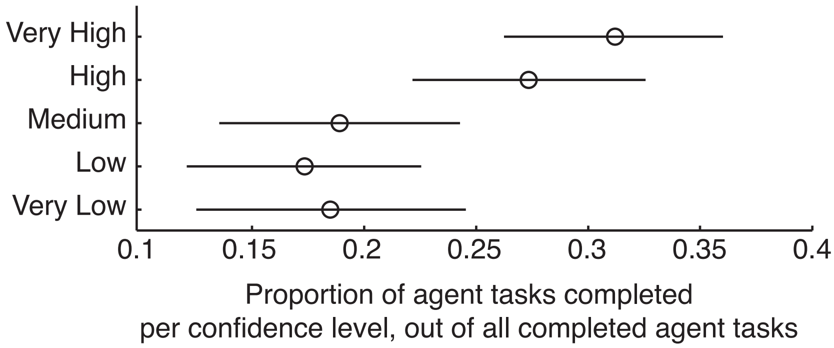

Figure 5.4: Means comparison for completed agent tasks per confidence level across all confidence levels, with the 95% confidence bars.

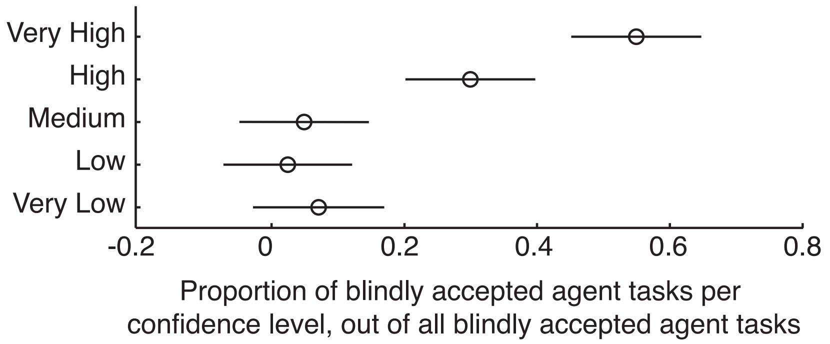

Figure 5.5: Means comparison for blindly accepted agent tasks per confidence level across all confidence levels, with the 95% confidence bars.

very low confidence level. Figure 5.5 shows the means comparison of the proportion of blindly accepted agent tasks per confidence level for all confidence levels.

Reward. A two-way ANOVA revealed a significant effect of incentive scheme ($p < 0.05$), but revealed no statistical significance across both displays of confidence information, with no interaction effects between the two. More reward were significantly earned in the *manual-incentive* ($M = 9.52, SD = 0.80$) than in the *agent-incentive* condition ($M = 8.65, SD = 1.62$).

Time. Participants took longer to complete agent tasks ($M = 49.48, SD = 24.84$) than manual tasks ($M = 35.22, SD = 11.29$) and a one-way ANOVA test indicates that this difference is significant ($p < 0.001$). Furthermore, a two-way ANOVA revealed a significant effect of incentive scheme ($p < 0.05$), but revealed no statistical significance across both displays of confidence information, with no effect of interaction between the two. A post-hoc Tukey test revealed that participants in the *agent-incentive* condition took significantly more time ($M = 49.60, SD = 11.62$) than participants in both the *no-incentive* ($M = 37.03, SD = 12.91$) and *manual-incentive* ($M = 37.07, SD = 14.13$) conditions.

Correct submissions. A two-way ANOVA revealed no statistical significance across incentive schemes and displays of confidence information, with no effect of interaction between the two.

Summary. In summary, the quantitative analysis of our data revealed that the display of *confidence* led participants to work on a higher proportion of agent tasks (top of Figure 5.3) and also blindly accept a higher proportion of agent tasks. Within the *confidence* condition, participants were more likely to work on tasks with *very high* confidence than any other tasks (Figure 5.4), and to blindly accept tasks with *very high* confidence more than tasks with *high* confidence, and these in turn more than tasks with lower levels of confidence (Figure 5.5). In terms of reward, the *agent-incentive* condition led participants to work on a higher proportion of agent tasks (bottom of Figure 5.3). In the *manual-incentive* condition participants gained a higher reward, while in the *agent-incentive* condition they spent more time on average per task.

5.3.2 Qualitative analysis

The next subsection details the results of our analysis of the interview data.

5.3.2.1 Selecting between manual and agent task

Our participants indicated a number of factors that they considered when choosing which type of task to complete. These factors included task easiness, engagement and also the reward associated with the task.

Choosing manual task. All but one participant in each of the *no-incentive* (19) and *manual-incentive* (19) conditions, and few participants (3) in the *agent-incentive* condition reported that one reason for preferring manual tasks was its easiness. Because manual tasks were written in a familiar language, they felt they were easier and quicker to complete than agent tasks. Furthermore, 14 participants in the *manual-incentive* condition also mentioned that the associated reward was a contributing factor for choosing manual tasks. Others were dismissive about the reward being a factor.

Choosing agent task. For half of the participants in the *agent-incentive* condition, easiness was a factor for preferring agent tasks. They reported that it was easier for them to compare textual snippets rather than completing a task that required grammatical reasoning. Two participants in the *agent-incentive* condition and 1 from each of the *no-incentive* and *manual-incentive* conditions mentioned that the challenging nature of agent tasks was also a contributing factor. Moreover, 14 participants in the *agent-incentive* condition indicated that the associated reward was also a key factor for preferring agent tasks. In contrast, 3 of participants in the *agent-incentive* condition were dismissive that the reward influenced their inclination to agent tasks. For example, P1, a 22-year-old male undergraduate studying Languages and in the *no-confidence,* *agent-incentive* condition, told us that *"Actually the money wasn't so much an issue. It was rather I wanted to test myself on something a bit harder"*.

There were also participants in the *agent-incentive* condition who reported that they were initially set to only do manual tasks since they felt more comfortable with them. However, they ended up also completing agent tasks. They reported realising that the manual task was harder than they thought. They also mentioned not consistently earning the reward through the manual tasks.

Switching between tasks. There were 44 participants (none in the *manual-incentive, no-confidence* condition) that shifted between the manual and the agent task, mostly to have a bit of variation. There were participants who mentioned switching between tasks when they were unsuccessful with one type of task. They felt that switching would help them in re-adjusting their focus.

In summary, some of our participants confirmed our findings from the statistical analysis that reward affected their task choice. Additionally, it appears that the associated reward also affected our participants' engagement and perception of task difficulty. In the next subsection, we describe how our participants used the confidence information to complete their tasks.

5.3.2.2 Utilising the confidence information

Interpretation. Our participants had various interpretations about the confidence levels. In the *confidence* condition, 17 reported that the agent tasks with *high* and *very high* confidence were very accurate and only the tasks with lower confidence levels have mistakes. In contrast, 7 participants in the *confidence* condition felt that the confidence information was not an important indicator of the amount of effort required for each confidence level. P2, a 22-year-old female undergraduate studying Languages and in the *manual-incentive, confidence* condition, argued that: *"...it [the agent] is not saying that it's right, it's saying its confidence but it's still not a 100% [sure]"*. Finally, 6 participants in the *confidence* condition disregarded agent tasks and completely focused on manual tasks, which in turn made them entirely disregard the confidence information.

Actual usage. We noticed from our observations that participants who interpreted the confidence display correctly devised different strategies in order to complete the tasks. This was confirmed through the interviews, as can be seen in the following.

Prioritise high confidence agent tasks. In the *confidence* condition, 11 participants focused on completing all high confidence agent tasks before either reviewing low confidence agent tasks or manual tasks. For example, P3, an 18-year-old male undergraduate studying History and in the *agent-incentive, confidence* condition, described how he utilised the confidence levels:

> *"The high and the very high I just clicked accept every time. I sort of found out that they seemed to be good enough to not warrant a review. Mediums I found out, anything below medium, I reviewed myself and it made sense for you to do ones that were, the one more confident first. So if you have loads of very lows because of the fact that you're only limited to 30 [submissions], I thought it was better to do the more confident ones first and leave the very low ones if you could. If a medium came, do that instead of a very low."*

Nine of these participants exclusively blindly accepted high confidence agent tasks because they felt they could rely on the agent to do a good job. For example P4, a 22-year-old female undergraduate studying Languages and in the *agent-incentive, confidence* condition, told us that:

> *"It felt more reassured if I clicked [review on] the low one, found errors then I think, OK I've done the job so it's probably more likely that I've found something and corrected it as opposed to if it was high, I would be doubting myself. I didn't find anything that could be wrong and it was just easier knowing there was definitely errors in the low ones that I could correct."*

While this participant initially blindly accepted *very high* confidence agent tasks, it later made her feel uneasy. As such, she decided to complete lower confidence agent tasks and manual tasks to make her feel that she is doing actual work. The other participants started reviewing high confidence tasks since they perceived those tasks would likely have the least amount of errors and therefore would require the least amount of effort.

Prioritise manual or low confidence agent tasks. In this strategy, 6 participants in the *confidence* condition focused on either reviewing manual tasks or low confidence agent tasks first before completing high confidence agent tasks. While all of these participants also believed that the high confidence agent tasks were reliable enough to not need reviewing, they intentionally left the high confidence agent tasks at the end. Four of these participants completed high confidence agent tasks towards the end of the study by blindly accepting them. The other two also intended to blindly accept high confidence tasks but managed to reach the submission limit. P5, a 21-year-old female undergraduate studying Medicine and in the *agent-incentive, confidence* condition, explained to us why she switched and focused on reviewing low confidence agent tasks:

> *"I switched it out to really low because high and very high were so accurate that it kinda got boring. And I thought if I'm gonna read through them I might as well find something to correct. And so I did that for quite a while. I was checking the time as well so once I knew I was sort of running out of time, I wanted to swap something where there was less [mistake] but still*

> *something to do, which is when I swapped to medium. And I decided to put it back to high and very high when I was sort of running out of time so that I knew I could scan them a lot quicker, get through as much as possible and then when I finally ran out of time I just clicked accept on all the high ones [laughing]."*

These statements demonstrate strategies adopted by our participants in utilising the confidence display. In contrast, the rest of the participants in the *confidence* condition completed their trials without reference to the confidence display. As a result, their approach to completing tasks was similar to the approach of participants in the *no-confidence* condition, which we will detail in the next subsection.

5.3.2.3 Making sense of the agent without confidence information

All the participants who reviewed agent tasks in the *no-confidence* condition and most (8 out of 13) of the participants in the *confidence* condition reported reviewing (rather than blindly accepting) agent tasks because they could not rely on the agent. For example, P6, a 36-year-old PhD student in Finance and in the *agent-incentive, confidence* condition told us that: *"I discovered that there's some mistake in there [agent task]. I don't think the software can read the paragraph"*. Others related their lack of reliance to personal experience with other software. P7, a 27-year-old female PhD student in Social Statistics and Demography and in the *no-incentive, no-confidence* condition, told us that: *"I've written in foreign languages before using just the computer into Word. When it changes, it gives you suggestion to change the grammar [and] it's not usually correct, [especially] if you are using a foreign language"*. The reliability of the agent was also a factor for participants to not complete agent tasks at all. P8, a 26-year-old PhD student in Business, in the *no-incentive, no-confidence* condition explained to us that:

> *"So if I was just to accept without reviewing, I would be fully trusting the software ... Since I didn't do it [agent task], I decided not to just fully trust the software then obviously going for something I could at least review [manual task]. Perhaps I was worse than the software [laughing] but at least it gave me a sense of control in the situation. Been able to carry out the task in my favour."*

In addition, unlike participants in the *confidence* condition who blindly accepted agent tasks because they were able to rely on the agent, the 8 participants in the *no-confidence* condition blindly accepted agent tasks only to gamble or try out their luck to earn money easily. P9, a 19-year-old female undergraduate studying English Literature and in the *manual-incentive, no-confidence* condition, mentioned that: *"Cause I just thought if you* *just chance it [press Accept], it would help boost it [the amount of reward] more quickly because just little 20p there there and it takes less time to click Accept"*. In fact, there were also participants that pressed the *Accept* button when they were running out of time as an attempt to earn as much money as possible. There were also 3 participants in the *no-confidence* condition blindly accepted agent tasks because they were bored and "*wanted to try the software*".

We also asked participants who completed agent tasks but never or very rarely blindly accepted agent tasks what changes can be made to increase their usage of blind acceptance. Participants mostly commented on changing the experimental setup, such as having unlimited submission counts, unlimited available tasks to complete with no involved penalty, reducing the amount of available time or increasing the tasks that needs to be completed could increase blind acceptance. Additionally, participants mentioned that accepting tasks blindly can be increased through knowing more about how the software works or knowing that the software is always correct. There were also some claims that regardless of any changes, they would not blindly accept tasks at all because it is their attitude towards software in general. For example, P10, a 23-year-old female MSc student in Finance and in the *manual-incentive, confidence* condition, told us that: *"I mean all those software because they are human made right so it might create errors and to choose between my error and the software error I'd rather choose mine"*.

Generally speaking, users prefer to rely on themselves as opposed to relying on the agent blindly when money is involved, especially when the confidence information of the agent is not displayed. In fact, blind acceptance without the confidence information is not perceived as trusting the agent, but more of taking chances to earn money easily. The next subsection will discuss issues reported related to the user interface.

5.3.2.4 UI issues

Participants also reported issues related to the UI of the interactive system. Firstly, 32 out of 60 participants explicitly mentioned that the sound alert distracted and stressed them and some even preferred the sound alert to be turned off. There were also two agent tasks that were accidentally blindly accepted and the participants who did this reported that they pressed the Accept button from the notification screen by mistake, while they wanted to press Review instead. Finally, two participants from the *confidence* condition also mentioned accidentally blindly accepting an agent task with low confidence, while they wanted to blindly accept a high confidence agent task instead.

5.4 Discussion

Here we discuss the results and the key lessons learnt from our study. Specifically, we discuss our experimental methodology and the effects of displaying agent confidence information.

5.4.1 Financial incentives and experimental method

The statistical analysis of our results revealed that the incentive scheme had an effect on the type of tasks that participants chose to complete. A higher proportion of agent tasks were completed in the *agent-incentive* condition than both the *no-incentive* and *manual-incentive* conditions. In other words, participants were sensitive to the financial incentives, and completed more of the type of tasks for which they received higher incentives. This result confirms our hypothesis H3, and validates the design of our study method, in that it demonstrates that the use of financial incentives was successful in motivating participants to do a specific task. In the *no-incentive* condition (where participants received a fixed £6 reward regardless of performance and task choice), participants completed more of the manual tasks, which is the one that requires the least effort. Users' sensitivity to financial incentives in the *agent-incentive* condition (i.e., more agent tasks get done) also indicates that participants are more inclined to use the agent when it provides higher utility than the manual task. In our study, experimental financial reward mimicked a situation in which the agent performs a task that is practically useful to participants (a real life example would be saving money on the energy bills by automatically controlling the thermostat). However, it should be noted as a limitation that the game-like nature of our experiment (including its limited duration) may have influenced participants to give more importance to the financial incentives than would be observed in real life. In other words, participants in the study may feel compelled to try and *win* as much as they can, just because it is a game (Deterding et al., 2011).

Even though our quantitative data clearly shows that the incentive scheme and the confidence information both had statistically significant effects on participants' behaviours, in the interviews participants suggested that a more complex and varied set of factors influenced the choice of tasks to complete. Most participants suggested that reward was only *one* contributing factor for preferring a task, while some went as far as completely dismissing the idea that the reward influenced their behaviour. Other reported factors included how easy or how challenging the task was perceived to be. At the same time, only participants in the *agent-incentive* condition described agent tasks as easier, and these are the tasks for which they received higher incentives. Furthermore, the general majority of participants reported manual tasks to be easier, and hence preferable. Therefore, the perception of a task as 'easy' seems to be influenced by the financial incentives.

It is possible that such bias was unconscious, or that participants felt embarrassment to acknowledge that they are driven by money. Such contrast between the quantitative results and the findings from the interviews reminds us that self-report may not always be dependable on its own, especially when attitudes towards financial incentives are involved.

In addition, the incentive scheme also had a statistically significant effect on the financial reward gained. Participants in the *manual-incentive* condition (where the manual task was rewarded more) earned more money than participants in the *agent-incentive* condition, suggesting that the manual task was easier than the agent task, as we intended. Such difference in effort required was further confirmed by another result of our analysis: participants took longer to complete agent tasks, on average, than to complete manual tasks.

5.4.2 Displaying the confidence information

The confidence information made a difference in how our participants interacted with the agent. We specifically hypothesised that there would be more agent tasks completed in the *confidence* condition than in the *no-confidence* condition (H1a). Our statistical analysis shows that a higher proportion of agent tasks were performed when the agent displayed the confidence information, confirming hypothesis H1a. In particular, participants completed more tasks with *very high* confidence level than tasks with lower levels of confidence, according to our hypothesis H1b. These results suggest that the different confidence information informed users about the amount of effort required before actually starting the tasks. Similarly, participants in the *confidence* condition blindly accepted a higher proportion of agent tasks, than in the *no-confidence* condition, confirming H2a. Furthermore, agent tasks with *very high* confidence were blindly accepted more than those with *high* confidence, and these in turn were blindly accepted more than tasks with lower confidence, confirming H2b.

The display of confidence information enabled participants to rely on the autonomous agent more. This result is in line with prior work on displaying confidence information (Beller et al., 2013; Helldin et al., 2013; McGuirl and Sarter, 2006). In turn, and as expected, our participants were unable to make an informed decision about using the agent when they had no confidence information. This result is also similar to findings from prior studies (Alan et al., 2014; Yang et al., 2014), even though our work is based on a different study method and different application (not energy related). To further support the quantitative data on this aspect, the interviews revealed a striking contrast between the *confidence* and *no-confidence* conditions. On the one hand, when confidence information was displayed most participants reported taking it into account for gauging their expectations about the performance of the agent, and in turn for choosing which tasks to perform. On the other hand, without confidence information available, participants described how they resorted to alternative ways to make sense of the agent, and to set their expectations. For example they referred to prior experience with software that they considered similar, like a spell checking software. However, such similarities may be based on superficial aspects of the agent, and hence be insubstantial, with the associated risk of generating incorrect expectations. To summarise, displaying the agent's confidence information allowed users to form strategies about how to utilise the agent based on their own attitude. Hence, the confidence information also increased the usage of the autonomous agent. We elaborate on these strategies in the next subsection.

5.4.3 Subjective perception and attitude

In general, our participants employed different strategies in utilising the confidence information, reflecting different personal attitudes toward autonomous agents. For example, some participants dismissed the confidence information, and the agent operation in general, based purely on their experience with other different computational agents. Other participants reported a preference for maintaining some form of control, similar to what has been reported in prior work (Alan et al., 2014). Others still acknowledged the meaning of the confidence information, but they favoured manual tasks, or agent tasks with lower confidence because they considered them more challenging, and hence rewarding. This finding aligns with the results of a study by (Ariely et al., 2008), who found that participants completing meaningful tasks were more motivated to work than participants who are working on less meaningful tasks. In our study, earning rewards by reviewing agent tasks was recognised by some participants as a more meaningful endeavour than simply earning rewards by blindly accepting agent tasks.

The interview data also revealed different user perceptions of the confidence information. Most participants were able to pick up on how well the confidence levels correlated to the reliability of the agent's output. These participants would describe that "*the chances were a lot higher*" for agent tasks with a *high* confidence level or higher to be correct, while they felt that "*there probably would be at least one mistake*" for agent tasks with a *medium* confidence level or lower. However, not all participants perceived that the confidence information related to the agent's capability. Some participants felt that they "*couldn't really distinguish a pattern*" and that the agent was only "*saying its confidence, it's still not a 100% positive*". This mindset of not relying on the confidence information emerged from participants who reported that they do not trust technology that can be considered similar to the one used in this study. It should be noted that only a minority of the participants (13) in the confidence condition reported such an attitude. Indeed the quantitative results indicate that, *in general*, displaying the confidence information makes a significant difference to the usage of autonomous agents. In summary, the user's perception of the display of confidence information is affected by the user's willingness to trust the agents that produce it. In our study, even though the confidence information provided was a reliable estimation of the correctness of the agent's output, there were still participants who disregarded it – a result of their reservations about trusting autonomous agents.

5.4.4 Reflecting on overall performance

No statistically significant effects of confidence information were found on the total reward gained by participants, nor on task completion time. The reward can be considered a proxy for the participant's overall performance in the experiment. While this finding is not conclusive (a larger sample size may reveal statistically significant differences), it does suggest that the confidence information did not influence overall performance. This result is perhaps counter-intuitive, because confidence led participants to blindly accept a higher proportion of (higher confidence) tasks, making them in principle more productive. Indeed, this result is in contrast with previous studies showing that displaying confidence information can improve user performance (Antifakos et al., 2004; McGuirl and Sarter, 2006; Rukzio et al., 2006). One possible explanation here is that the time gained by blindly accepting tasks was spent in an unproductive way (unproductive in terms of the experiment's financial reward). Indeed the interviews suggest that some participants preferred tasks that are more challenging, rather than easier, or tasks for which they have more control, because they were generally sceptical about the agents' abilities and disregarded the confidence information.

5.5 Implications

A number of implications follow from this analysis, which we portray in the following:

- First, it appears that displaying confidence information can increase the uptake of autonomous agents. Confidence information provides users an indication about how much effort and attention is required from them. However, our results focus on the initial acceptance of autonomous agents. Future work should focus on whether displaying confidence information can also help incentivise users to maintain interaction with autonomous agents. In addition, future work should also investigate whether different numbers of confidence levels also have an effect on how users utilise the confidence information.
- Second, the different strategies that our participants employed to utilise the confidence information suggested that users have different preferences about how they would adopt autonomous agents. Without much knowledge of how exactly the agent works, users will likely interpret the confidence levels based on their initial

experience of the agent itself or experience with different technologies. While users could be informed about how to deal with different confidence levels by explicitly suggesting at what points should the agent be relied on and when users should intervene, users may eventually interact with the agent differently as their experience using it increases. Combining confidence information with other techniques may improve its utilisation, such as providing short explanations as to why an agent-completed task is of high confidence (Lim et al., 2009) or why the agent might make mistakes (Dzindolet, 2003).

- Third, participants highlighted weaknesses in the prototype used in our study, leading to some practical suggestions for the design of future interfaces for agent-based systems. For example, the sound alert was found to be distracting and stressful by most of our participants. This result aligns with the study of (Parasuraman and Miller, 2004), where it has been found that good etiquette in terms of the agent's communication style (e.g. notifying user to complete a task only when the user is not currently engaged with another task) can increase user's reliance on the agent and also improved their performance in diagnosing an agent's mistakes. More basic usability issues include verifying users' action to avoid accidental review or blind acceptance of agent tasks.

- Finally, our study highlights the effectiveness of using behavioural economics research methodologies to study user interactions with agent-based systems Further studies should look into applying this methodology to investigate acceptance of other technologies, such as recommender systems. Moreover, studies should look into the validity of this approach for other types of users, such as workers in crowdsourcing marketplaces [4], who may react differently to various financial incentives.

5.6 Limitations

We opted to not inform participants that the OCR agent was purely simulated. As mentioned in Section 5.2.1, this was done to ensure consistency since all participants encounter the same level of quality from the agent. While some participants may have thought that the agent was not simulated, others may not think that way, which may have potentially polluted the data. However, we note that during the interviews, none of the participants thought that the OCR agent was not real, even with participants with a technical background. Once they were debriefed, none of our participants felt that their behaviour would have changed had they known. Despite this, future studies should investigate interactions with actual software agents in a real-world applications, to guarantee that participants will act in a realistic manner. In addition, while our study method provides higher ecological validity than traditional lab setting, future work

[4] http://www.mturk.com/

should reconfirm our findings in a real world setting to validate the effectiveness of our study method. Furthermore, even though our method places considerable emphasis on financial rewards as a motivational factor for using (or ignoring) the autonomous agent, the interviews revealed that a variety of other factors are also at play (e.g. curiosity, challenge, etc.). Although our method has proven to be flexible enough to allow these factors to emerge, future work should investigate situations where financial effects are not in the picture at all.

Regarding the confidence information, this work relies on the availability of accurate confidence information, such as when the agent uses an appropriate model to learn the data. However, it is important to point out that this may not always be the case. Further research is needed to evaluate the effects of unreliable confidence information. Lastly, future studies should investigate longer term effects and also how people would react to finer- or coarser-grained confidence levels.

5.7 Summary

Previous studies have tested different feedback mechanisms to help people interact with autonomous agents. In particular, studies have looked at how to increase the transparency of agents, such as by providing information about each step of of an agent or explanations about why an agent did certain actions. Although these strategies were found to be helpful, they often overload the user and at times, they can also be difficult for the user to understand. In contrast, other studies have explored whether displaying an agent's confidence information can improve people's interaction with it. Extending prior work, we performed a lab study with 60 participants to test whether confidence information can increase the adoption of an agent and help people correctly delegate agency to it, in a situation where the agent can be ignored. Such a situation mimics real-life situation, where users have the freedom to use whether they use an agent or not. A combination of quantitative and qualitative data revealed that when confidence information is available users are more likely to take advantage of the agent. This result can be explained through the observation that users can be guided in selecting which agent tasks to concentrate on by displaying the confidence information.

An important implication of our work, then, is that if at all possible, confidence information should be included in the feedback from autonomous agents. This then addresses our research question R3, suggesting that displaying an agent's confidence information is an effective feedback mechanism that can help people to delegate agency to an autonomous agent, such that people do not rely too much or too little on them, especially when there are uncertainties in the performance of the agent. Moreover, through a comparison of the effects of different incentive schemes our study also demonstrates that our participants were sensitive to different reward mechanisms. Such findings further suggest that our study design, similar to the group buying study, is effective to realistically evaluate interactions with autonomous agents in a controlled lab setting.

Chapter 6

Conclusions and Future Work

This chapter gives a detailed summary of this thesis and also provides potential avenues for future research.

6.1 Conclusions

In Chapter 1, we stated that while autonomous agents are becoming more prevalent, uncertainty about their capabilities and their quality of work has resulted in people not being able to effectively use them. We highlighted that there is little knowledge about designing interactions that will enable people to effectively use autonomous agents in *non-specialist applications*, applications where users are not trained experts of the systems. Based on this premise, we mentioned that the main objective of this thesis is to provide design guidelines to help people deal with the uncertainties of using autonomous agents. We concluded that by achieving this goal, user adoption and utilisation of autonomous agents can be improved.

In Chapter 2, we reviewed previous studies of how people interact with autonomous agents in various applications, such as in industrial use and domestic applications. Our review revealed the lack of understanding about factors that enable agency delegation, particularly in situations where the agents being used are subject to uncertainties. We also highlighted that while studies suggest that people are willing to adopt autonomous agents, findings also showed that users do not effectively make use of them, as they often have difficulty in granting and revoking agency. We also learnt that most studies related to users interacting with autonomous agents lack ecological validity and no research so far has looked at agents supporting food-related practices. Finally, we also provided an overview of behavioural economics literature, inspiring the studies we conducted.

In Chapter 3, we reported a qualitative study with 11 households, designed to observe how people deal with uncertainty through the veg box scheme, as an example of an unpredictable service that exemplifies agency delegation. Through semi-structured interviews and 14 days of diary study, our findings revealed how participants integrated the veg box in their food-related practices, such as in food shopping and preparation. Furthermore, and perhaps more importantly, our study extends prior work as it suggests that agency delegation can also be facilitated by factors that go beyond efficiency and convenience. This includes factors such as its ability to enable creativity through the 'element of surprise' and its ability to support consumers' personal values. We provided key design implications for future autonomous agents, which includes supporting user creativity, allowing users to express their values, and being able to coordinate an agent's actions with the users' own coordinated practices. For example, cooking robots such as Moley[1] should be designed to take into account creativity, by exploring various recipes that the user may not have tried before, but uses the same ingredients the user is accustomed to using.

In Chapter 4, we further investigated the importance of incorporating personal human values in the design of autonomous agents. Going beyond the previous chapter, we moved to a more concrete perspective by evaluating how users interact with an actual software agent. Specifically, we tested whether people would be more inclined to adopt a software agent designed with human values or a software agent that is only goal-oriented, whereby there are uncertainties in the performance of both agents. In particular, we presented a study designed to explore how *fairness*, a well-known human value in bargaining situations, affects the adoption of autonomous group buying shopping agents. In more detail, 20 participants were exposed to two alternative group buying autonomous shopping agents: one designed around fairness, and the other is goal-oriented. Results gathered through quantitative and qualitative methods showed that while all participants welcomed the support provided by the autonomous agents, participants were more accepting of the agent that takes fairness into account. Such implication can be applied to future shopping assistants or agents for bargaining, whereby such technology should not only focus on the goal of finding a group, but it should also ensure that the user is getting a fair deal. Moreover, other than incorporating human values in agent design, our other key design implication is to ensure that users are aware of the agent's level of service to avoid expectation mismatch, so as to help users handle uncertainties of an agent's actions.

In Chapter 5, we further investigated the effectiveness of displaying an agent's level of service in improving user adoption of the technology. Specifically, in this chapter we presented a lab study with 60 participants, designed to investigate whether users can make sense of confidence information and whether it can influence users' attitude towards autonomous agents. A combination of quantitative and qualitative data revealed that when confidence information is available users are more likely to take advantage of the agent, compared to a situation where such information is not available. This result can

[1] http://www.moley.com/

be explained through the observation that users can be guided to complete tasks that involves interacting with an agent by displaying the confidence information. The main implication of this chapter, then, is that if at all possible confidence information should be included in the feedback from autonomous agents, not only to increase the chances of their uptake, but to also help users know when to let the agent do its work autonomously and when to manually take over. For instance, autonomous technology such as the Nest Thermostat should include confidence information to provide users some estimation of how well the sensing works and whether it needs to be checked or not.

In addition, the design of the lab studies described in both Chapters 4 and 5 were inspired by behavioural economics methods as described in Section 2.3.1. In more detail, both studies employed multitasking and financial incentives to incorporate *opportunity cost* as a way of simulating realism. Doing so enabled us to realistically evaluate interactions with autonomous agents in a controlled lab setting.

To summarise, the studies reported in this thesis have provided a better understanding of how people deal with the uncertainties of autonomous agents in non-specialist applications. Our design implications serve guidance on how to design interactions that will support effective agency delegation and improve the adoption of agents. Another contribution of our work is the study method we used to investigate user interactions with autonomous agents in a lab setting. Indeed, the outcomes of this research have tackled the research challenges and addressed the research questions we set out to achieve. Nonetheless, we acknowledge that our studies have limitations. We believe such caveats can lead to potential future research, which we discuss in the next section.

6.2 Future Work

The studies detailed in this report highlighted a number of future research opportunities. We discussed limitations and avenues for future work that are specific to each study in their respective chapters. In this section, we will discuss potential future work in light of the thesis as a whole.

Both the veg box and group buying study have shown that people are willing to make use of autonomous agents that support their food-based practices. In particular, our studies look at the feasibility of agents that performs food shopping on behalf of the users. However, our work only scratched the surface of the domain of food, especially since we only focused on activities related to food shopping. Future studies should explore interactions with autonomous agents supporting other food-related activities, such as food preparation, cooking and disposal. Furthermore, future studies should also investigate the adoption of *physical* autonomous agents. For example, in one of our studies (excluded in this thesis), we explored how people's perception of self-cleaning robots can be changed through motion (Garcia et al., 2016). In more detail, participants who saw a vacuum robot docking as it finishes its task thought that it performed better than a vacuum robot that has already docked to its station. This study can be integrated in future versions of the work stated in this thesis by examining whether motion can also affect a person's inclination to delegate agency to autonomous agents.

There are also avenues for future research related to the study designs used in our work. For instance, we designed our lab studies to include higher ecological validity than traditional controlled lab settings. However, it is worth noting that users may still react differently in a real-world situation, as many other factors will likely come to play. In addition, the duration of our studies, especially those conducted in the lab, were limited to short periods of times. Future studies should examine whether the findings we observed hold or whether people's behaviours change when interacting with agents for longer periods of time. Moreover, although people who took part in our studies were not all students, almost all of them were educated to above average levels. As as result, our participants do not fully represent the overall society and therefore our work should be extended to a more general population. Addressing these proposed changes in the methods of our studies can further strengthen the claims stated in this thesis.

In connection to the above, there are also limitations regarding the agents we used in all of our studies, in that they lacked realism. In the veg box study, the 'agent' was in the form of a service formed by a group of people. In the group buying study, the agents were designed to fit in a market simulation of how we believed a group buying market would work in real life. In the confidence information study, the agent used were based on a Wizard-of-Oz approach and hence was also not a real working software. While these choices were intentionally made for the purpose of our studies, future work should endeavour to examine interactions with real autonomous software agents. Such agents will behave more realistically and more unpredictable to users, making the observations between the users and the agents even more meaningful.

Finally, future work should investigate other areas that involves the use of autonomous agents in everyday life. For instance, machine learning is becoming increasingly embedded in autonomous agents. As such, research should focus on helping non-technical users utilise agents with machine learning capabilities, especially those that are liable to failure. Autonomous agents are also becoming more prevalent in IoT devices, especially in the home environment. Since IoT devices at home are likely to be used by multiple users at the same time, research should also investigate how to design interaction mechanisms that cater for multiple users interacting with their IoT devices. Moreover, blockchain technology, which allows data to be stored in an encrypted distributed ledger format, may be used to track the performance of autonomous agents. This paves way for future research to examine how autonomous agents that employs blockchains should be designed, especially in helping non-technical users make sense of these technologies. Indeed, we hope that the work presented in this thesis will stimulate discussions and future work around interaction design for autonomous agents in everyday life.

References

U. G. C. S. Adviser. Distributed ledger technology: Beyond block chain. Technical report, UK Government Office for Science, January 2016.

A. Alan, E. Costanza, J. Fischer, S. D. Ramchurn, T. Rodden, and N. R. Jennings. A field study of human-agent interaction for electricity tariff switching. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 965–972, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

A. T. Alan, E. Costanza, S. D. Ramchurn, J. Fischer, T. Rodden, and N. R. Jennings. Tariff agent: Interacting with a future smart energy system at home. *ACM Trans. Comput.-Hum. Interact.*, 23(4):25:1–25:28, August 2016a.

A. T. Alan, M. Shann, E. Costanza, S. D. Ramchurn, and S. Seuken. It is too hot: An in-situ study of three designs for heating. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 5262–5273, New York, NY, USA, 2016b. ACM.

R. Alijani, M. Farhadi, M. Ghodsi, M. Seddighin, and A. S. Tajik. Envy-free mechanisms with minimum number of cuts. In *AAAI*, pages 312–318, 2017.

S. Antifakos, N. Kern, B. Schiele, and A. Schwaninger. Towards improving trust in context-aware systems by displaying system confidence. In *Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices &Amp; Services*, MobileHCI '05, pages 9–14, New York, NY, USA, 2005. ACM.

S. Antifakos, A. Schwaninger, and B. Schiele. Evaluating the effects of displaying uncertainty in context-aware applications. In N. Davies, E. Mynatt, and I. Siio, editors, *UbiComp 2004: Ubiquitous Computing*, volume 3205 of *Lecture Notes in Computer Science*, pages 54–69. Springer Berlin Heidelberg, 2004.

D. Ariely, E. Kamenica, and D. Prelec. Man's search for meaning: The case of legos. *Journal of Economic Behavior & Organization*, 67(34):671 – 677, 2008.

F. Asselin and B. Chaib-draa. Toward a protocol for the formation of coalitions of buyers. In *Proceedings of the 5th International Conference on Electronic Commerce Research (ICECR-5)*, 2002.

M. Ball and V. Callaghan. Introducing intelligent environments, agents and autonomy to users. In *Intelligent Environments (IE), 2011 7th International Conference on*, pages 382–385, July 2011.

C. Basu, Q. Yang, D. Hungerman, M. Singhal, and A. D. Dragan. Do you want your autonomous car to drive like you? In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '17, pages 417–425, New York, NY, USA, 2017. ACM.

J. Beller, M. Heesen, and M. Vollrath. Improving the driver-automation interaction: An approach using automation uncertainty. *Human Factors*, 55(6):1130–1141, 2013.

T. Bosona, G. Gebresenbet, I. Nordmark, D. Ljungberg, et al. Box-scheme based delivery system of locally produced organic food: Evaluation of logistics performance. *Journal of Service Science and management*, 4(03):357, 2011.

J. Bourgeois, J. van der Linden, G. Kortuem, B. A. Price, and C. Rimmer. Conversations with my washing machine: An in-the-wild study of demand shifting with self-generated energy. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '14, pages 459–470, New York, NY, USA, 2014. ACM.

V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.

B. Brown and E. Laurier. The trouble with autopilots: Assisted and autonomous driving on the social road. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 416–429, New York, NY, USA, 2017. ACM.

E. Brown, S. Dury, and M. Holdsworth. Motivations of consumers that use local, organic fruit and vegetable box schemes in central england and southern france. *Appetite*, 53 (2):183–188, 2009.

N. Buchan, R. Croson, E. Johnson, and G. Wu. *Gain and Loss Ultimatums*, pages 1–23. 2005.

C. Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2003.

C. F. Camerer, E. Johnson, T. Rymon, and S. Sen. Cognition and framing in sequential bargaining for gains and losses. *Frontiers of game theory*, 104:27–47, 1993.

M. E. Cecchinato, A. Sellen, M. Shokouhi, and G. Smyth. Finding email in a multi-account, multi-device world. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 1200–1210, New York, NY, USA, 2016. ACM.

E. Cha, J. Forlizzi, and S. S. Srinivasa. Robots in the home: Qualitative and quantitative insights into kitchen organization. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, pages 319–326, New York, NY, USA, 2015. ACM.

H. Chan and J. Chen. Budget feasible mechanisms for dealers. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, AAMAS '16, pages 113–122, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.

J.-H. Chen, P. P.-Y. Chi, H.-H. Chu, C. C.-H. Chen, and P. Huang. A smart kitchen for nutrition-aware cooking. *IEEE Pervasive Computing*, 9(4):58–65, October 2010.

P.-Y. P. Chi, J.-H. Chen, H.-H. Chu, and J.-L. Lo. Enabling calorie-aware cooking in a smart kitchen. In *Proceedings of the 3rd International Conference on Persuasive Technology*, PERSUASIVE '08, pages 116–127, Berlin, Heidelberg, 2008. Springer-Verlag.

N. Cila, I. Smit, E. Giaccardi, and B. Kröse. Products as agents: Metaphors for designing the products of the iot age. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 448–459, New York, NY, USA, 2017. ACM.

A. K. Clear, M. Hazas, J. Morley, A. Friday, and O. Bates. Domestic food and sustainable design: A study of university student cooking and its impacts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 2447–2456, New York, NY, USA, 2013. ACM.

Y. J. Cohler, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Optimal envy-free cake cutting. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, AAAI'11, pages 626–631. AAAI Press, 2011.

R. Comber, J. Hoonhout, A. van Halteren, P. Moynihan, and P. Olivier. Food practices as situated action: Exploring and designing for everyday food practices with households. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 2457–2466, New York, NY, USA, 2013. ACM.

E. Costanza, J. E. Fischer, J. A. Colley, T. Rodden, S. D. Ramchurn, and N. R. Jennings. Doing the laundry with agents: A field trial of a future smart energy system in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 813–822, New York, NY, USA, 2014. ACM.

H. Davis, B. Nansen, F. Vetere, T. Robertson, M. Brereton, J. Durick, and K. Vaisutis. Homemade cookbooks: A recipe for sharing. In *Proceedings of the 2014 Conference on Designing Interactive Systems*, DIS '14, pages 73–82, New York, NY, USA, 2014. ACM.

C. M. de Melo, S. Marsella, and J. Gratch. "do as i say, not as i do": Challenges in delegating decisions to automated agents. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, AAMAS '16, pages 949–956, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.

C. M. de Melo, S. Marsella, and J. Gratch. Increasing fairness by delegating decisions to autonomous agents. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, pages 419–425, Richland, SC, 2017a. International Foundation for Autonomous Agents and Multiagent Systems.

C. M. de Melo, S. Marsella, and J. Gratch. Social decisions and fairness change when people's interests are represented by autonomous agents. *Autonomous Agents and Multi-Agent Systems*, Jul 2017b.

E. G. de Tjerk, F. A. Henryk, and M. A. Neerincx. Adaptive automation based on an object-oriented task model: Implementation and evaluation in a realistic c2 environment. *Journal of Cognitive Engineering and Decision Making*, 4(2):152–182, 2010.

D. Dearman, A. Varshavsky, E. De Lara, and K. N. Truong. An exploration of location error estimation. In *Proceedings of the 9th International Conference on Ubiquitous Computing*, UbiComp '07, pages 181–198, Berlin, Heidelberg, 2007. Springer-Verlag.

F. Dehais, V. Peysakhovich, S. Scannella, J. Fongue, and T. Gateau. "automation surprise" in aviation: Real-time solutions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 2525–2534, New York, NY, USA, 2015. ACM.

M. Desai, P. Kaniarasu, M. Medvedev, A. Steinfeld, and H. Yanco. Impact of robot failures and feedback on real-time trust. In *Proceedings of the 8th ACM/IEEE International Conference on Human-robot Interaction*, HRI '13, pages 251–258, Piscataway, NJ, USA, 2013. IEEE Press.

S. Deterding, D. Dixon, R. Khaled, and L. Nacke. From game design elements to gamefulness: Defining "gamification". In *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments*, MindTrek '11, pages 9–15, New York, NY, USA, 2011. ACM.

R. R. Dholakia. Going shopping: key determinants of shopping behaviors and motivations. *International Journal of Retail & Distribution Management*, 27(4):154–165, 1999.

M. Dzindolet. The role of trust in automation reliance. *International Journal of Human-Computer Studies*, 58(6):697–718, June 2003.

M. Fan and K. N. Truong. Soqr: Sonically quantifying the content level inside containers. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '15, pages 3–14, New York, NY, USA, 2015. ACM.

U. Farooq and J. Grudin. Human-computer integration. *interactions*, 23(6):26–32, October 2016.

H. S. Ferdous, B. Ploderer, H. Davis, F. Vetere, K. O'Hara, G. Farr-Wharton, and R. Comber. Tabletalk: Integrating personal devices and content for commensal experiences at the family dinner table. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '16, pages 132–143, New York, NY, USA, 2016. ACM.

J. Fink. *Anthropomorphism and Human Likeness in the Design of Robots and Human-Robot Interaction*, pages 199–208. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

J. Forlizzi and C. DiSalvo. Service robots in the domestic environment: A study of the roomba vacuum in the home. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-robot Interaction*, HRI '06, pages 258–265, New York, NY, USA, 2006. ACM.

B. Friedman. Agents of value. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 1–2, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems.

B. Friedman and P. H. Kahn. Human agency and responsible computing: Implications for computer system design. *Journal of Systems and Software*, 17(1):7 – 14, 1992. Computer Ethics.

B. Friedman, P. H. Kahn, and A. Borning. Value sensitive design and information systems. *The handbook of information and computer ethics*, pages 69–101, 2008.

Y. Gal, S. Kraus, M. Gelfand, H. Khashan, and E. Salmon. An adaptive agent for negotiating with people in different cultures. *ACM Trans. Intell. Syst. Technol.*, 3(1): 8:1–8:24, October 2011.

Y. K. Gal, M. Mash, A. D. Procaccia, and Y. Zick. Which is the fairest (rent division) of them all? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 67–84, New York, NY, USA, 2016. ACM.

E. Ganglbauer, G. Fitzpatrick, and R. Comber. Negotiating food waste: Using a practice lens to inform design. *ACM Trans. Comput.-Hum. Interact.*, 20(2):11:1–11:25, May 2013.

E. Ganglbauer, G. Fitzpatrick, and F. Güldenpfennig. Why and what did we throw out?: Probing on reflection through the food waste diary. In *Proceedings of the 33rd*

Annual ACM Conference on Human Factors in Computing Systems, CHI '15, pages 1105–1114, New York, NY, USA, 2015. ACM.

E. Ganglbauer, G. Fitzpatrick, O. Subasi, and F. Güldenpfennig. Think globally, act locally: A case study of a free food sharing community and social networking. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '14, pages 911–921, New York, NY, USA, 2014. ACM.

P. G. Garcia, E. Costanza, S. D. Ramchurn, and J. K. M. Verame. The potential of physical motion cues: Changing people's perception of robots' performance. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '16, pages 510–518, New York, NY, USA, 2016. ACM.

W. W. Gaver, J. Beaver, and S. Benford. Ambiguity as a resource for design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, pages 233–240, New York, NY, USA, 2003. ACM.

U. Gneezy, E. Haruvy, and A. E. Roth. Bargaining under a deadline: evidence from the reverse ultimatum game. *Games and Economic Behavior*, 45(2):347 – 368, 2003. Special Issue in Honor of Robert W. Rosenthal.

M. A. Goodrich, T. W. McLain, J. D. Anderson, J. Sun, and J. W. Crandall. Managing autonomy in robot teams: Observations from four experiments. In *Proceedings of the ACM/IEEE International Conference on Human-robot Interaction*, HRI '07, pages 25–32, New York, NY, USA, 2007. ACM.

A. Grimes and R. Harper. Celebratory technology: New directions for food research in hci. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 467–476, New York, NY, USA, 2008. ACM.

W. Guth, R. Schmittberger, and B. Schwarze. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, 3(4):367 – 388, 1982.

R. Hamada, J. Okabe, I. Ide, S. Satoh, S. Sakai, and H. Tanaka. Cooking navi: Assistant for daily cooking in kitchen. In *Proceedings of the 13th Annual ACM International Conference on Multimedia*, MULTIMEDIA '05, pages 371–374, New York, NY, USA, 2005. ACM.

T. Helldin, G. Falkman, M. Riveiro, and S. Davidsson. Presenting system uncertainty in automotive uis for supporting trust calibration in autonomous driving. In *Proceedings of the 5th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, AutomotiveUI '13, pages 210–217, New York, NY, USA, 2013. ACM.

T. Helldin, U. Ohlander, G. Falkman, and M. Riveiro. Transparency of automated combat classification. In *Engineering Psychology and Cognitive Ergonomics*, pages 22–33. Springer, 2014.

E. Horvitz and J. Krumm. Some help on the way: Opportunistic routing under uncertainty. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, pages 371–380, New York, NY, USA, 2012. ACM.

A. Hupfeld and T. Rodden. Laying the table for hci: Uncovering ecologies of domestic food consumption. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pages 119–128, New York, NY, USA, 2012. ACM.

T. Ito, H. Ochi, and T. Shintani. A group-buy protocol based on coalition formation for agent-mediated e-commerce. *IJCIS*, 3(1):11–20, 2002.

N. R. Jennings, L. Moreau, D. Nicholson, S. Ramchurn, S. Roberts, T. Rodden, and A. Rogers. Human-agent collectives. *Commun. ACM*, 57(12):80–88, November 2014.

W. Jiang, J. E. Fischer, C. Greenhalgh, S. D. Ramchurn, F. Wu, N. R. Jennings, and T. Rodden. Social implications of agent-based planning support for human teams. In *Collaboration Technologies and Systems (CTS), 2014 International Conference on*, pages 310–317. IEEE, 2014.

D. Kahneman. *Thinking, fast and slow*. Macmillan, 2011.

D. Kahneman and A. Tversky. Choices, values, and frames. *American Psychologist*, 39 (4):341–350, April 1984.

R. J. Kauffman, H. Lai, and C.-T. Ho. Incentive mechanisms, fairness and participation in online group-buying auctions. *Electronic Commerce Research and Applications*, 9 (3):249–262, 2010.

R. J. Kauffman and B. Wang. Bid together, buy together: On the efficacy of group-buying business models in internet-based selling. In *Handbook of Electronic Commerce in Business and Society, BOCA*, pages 27–28. CRC Press, 2002.

G. Klien, D. D. Woods, J. M. Bradshaw, R. R. Hoffman, and P. J. Feltovich. Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19(6):91–95, Nov 2004.

S. Kuznetsov, C. J. Santana, and E. Long. Everyday food science as a design space for community literacy and habitual sustainable practice. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 1786–1797, New York, NY, USA, 2016. ACM.

A. Lazar, C. Koehler, J. Tanenbaum, and D. H. Nguyen. Why we use and abandon smart devices. In *Proceedings of the 2015 ACM International Joint Conference on*

Pervasive and Ubiquitous Computing, UbiComp '15, pages 635–646, New York, NY, USA, 2015. ACM.

J. D. Lee and K. A. See. Trust in automation: Designing for appropriate reliance. *Human factors*, 46(1):50–80, 2004.

K. J. Lee, Y. K. Joo, and C. Nass. Partially intelligent automobiles and driving experience at the moment of system transition. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 3631–3634, New York, NY, USA, 2014. ACM.

M. K. Lee, S. Kiesler, and J. Forlizzi. Mining behavioral economics to design persuasive technology for healthy choices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 325–334, New York, NY, USA, 2011. ACM.

Y. K. Lee, S. W. Kim, M. J. Ko, and S. E. Park. *Pricing Agents for a Group Buying System*, pages 693–700. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.

H. Lemelson, T. King, and W. Effelsberg. A study on user acceptance of error visualization techniques. In *Proceedings of the 5th Annual International Conference on Mobile and Ubiquitous Systems: Computing, Networking, and Services*, Mobiquitous '08, pages 53:1–53:6, ICST, Brussels, Belgium, Belgium, 2008. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

M. Lewis, H. Wang, S.-Y. Chien, P. Scerri, P. Velagapudi, K. Sycara, and B. Kane. Teams organization and performance in multi-human/multi-robot teams. In *Systems Man and Cybernetics (SMC), 2010 IEEE International Conference on*, pages 1617–1623, Oct 2010.

B. Y. Lim and A. K. Dey. Investigating intelligibility for uncertain context-aware applications. In *Proceedings of the 13th International Conference on Ubiquitous Computing*, UbiComp '11, pages 415–424, New York, NY, USA, 2011. ACM.

B. Y. Lim, A. K. Dey, and D. Avrahami. Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, pages 2119–2128, New York, NY, USA, 2009. ACM.

N. Lim. Consumers perceived risk: sources versus consequences. *Electronic Commerce Research and Applications*, 2(3):216 – 228, 2003. Selected Papers from the Pacific Asia Conference on Information Systems.

L. Lin and M. A. Goodrich. Sliding autonomy for uav path-planning: Adding new dimensions to autonomy management. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 1615–1624, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.

S. Liu, S. Helfenstein, and A. Wahlstedt. Social psychology of persuasion applied to human agent interaction. *Human Technology: An Interdisciplinary Journal on Humans in ICT Environments*, 2008.

J. M. McGuirl and N. B. Sarter. Supporting trust calibration and the effective use of decision aids by presenting dynamic system confidence information. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 48(4):656–665, 2006.

S. Mennicken, O. Zihler, F. Juldaschewa, V. Molnar, D. Aggeler, and E. M. Huang. "it's like living with a friendly stranger": Perceptions of personality traits in a smart home. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '16, pages 120–131, New York, NY, USA, 2016. ACM.

A. E. Milewski and S. H. Lewis. Delegating to software agents. *International Journal of Human-Computer Studies*, 46(4):485 – 500, 1997.

S. Moran, N. Pantidi, K. Bachour, J. E. Fischer, M. Flintham, T. Rodden, S. Evans, and S. Johnson. Team reactions to voiced agent instructions in a pervasive game. In *Proceedings of the 2013 international conference on Intelligent user interfaces*, pages 371–382. ACM, 2013.

L. Moreau and P. Groth. Provenance: An introduction to prov. *Synthesis Lectures on the Semantic Web: Theory and Technology*, 3(4):1–129, 2013.

C. Nass, Y. Moon, and P. Carney. Are people polite to computers? responses to computer-based interviewing systems. *Journal of Applied Social Psychology*, 29(5): 1093–1109, 1999.

T. Neumann, S. Schosser, and B. Vogt. Ultimatum bargaining over losses and gains an experimental comparison. *Social Science Research*, 67:49 – 58, 2017.

C. Neustaedter, A. J. B. Brush, and S. Greenberg. The calendar is crucial: Coordination and awareness through the family calendar. *ACM Trans. Comput.-Hum. Interact.*, 16 (1):6:1–6:48, April 2009.

D. A. Norman. The 'problem' with automation: inappropriate feedback and interaction, not 'over-automation'. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 327(1241):585–593, 1990.

D. A. Norman. How might people interact with agents. *Communications of the ACM*, 37(7):68–71, 1994.

J. Palay and M. Newman. Suchef: An in-kitchen display to assist with "everyday" cooking. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '09, pages 3973–3978, New York, NY, USA, 2009. ACM.

R. Parasuraman and D. H. Manzey. Complacency and Bias in Human Use of Automation: An Attentional Integration. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 52(3):381–410, June 2010.

R. Parasuraman and C. A. Miller. Trust and etiquette in high-criticality automated systems. *Communications of the ACM*, 47(4):51–55, 2004.

R. Parasuraman, R. Molloy, and I. L. Singh. Performance consequences of automation-induced 'complacency'. *The International Journal of Aviation Psychology*, 3(1):1–23, 1993.

C. Petrini. *Slow Food Nation: Why Our Food Should be Good, Clean, and Fair*. Random House Incorporated, 2007.

P. K. K. P.K. Kannan. Dynamic pricing on the internet: Importance and implications for consumer behavior. *International Journal of Electronic Commerce*, 5(3):63–83, 2001.

L. Pschetz, E. Tallyn, R. Gianni, and C. Speed. Bitbarista: Exploring perceptions of data transactions in the internet of things. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 2964–2975, New York, NY, USA, 2017. ACM.

S. D. Ramchurn, J. E. Fischer, Y. Ikuno, F. Wu, J. Flann, and A. Waldock. A study of human-agent collaboration for multi-uav. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, pages 1184–1192. AAAI Press, 2015.

B. Reeves and C. Nass. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press, New York, NY, USA, 1996.

W. H. Reitberger, W. Spreicer, and G. Fitzpatrick. Nutriflect: Reflecting collective shopping behavior and nutrition. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems*, CHI '14, pages 3309–3318, New York, NY, USA, 2014. ACM.

T. A. Rodden, J. E. Fischer, N. Pantidi, K. Bachour, and S. Moran. At home with agents: Exploring attitudes towards future smart energy infrastructures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 1173–1182, New York, NY, USA, 2013. ACM.

C. Rödel, S. Stadler, A. Meschtscherjakov, and M. Tscheligi. Towards autonomous cars: The effect of autonomy levels on acceptance and user experience. In *Proceedings of the*

6th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI '14, pages 11:1–11:8, New York, NY, USA, 2014. ACM.

Y. Rogers. Moving on from weiser's vision of calm computing: Engaging ubicomp experiences. In *Proceedings of the 8th International Conference on Ubiquitous Computing*, UbiComp'06, pages 404–421, Berlin, Heidelberg, 2006. Springer-Verlag.

M. Roos and J. Rothe. Complexity of social welfare optimization in multiagent resource allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS '10, pages 641–648, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems.

A. E. Roth, V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir. Bargaining and market behavior in jerusalem, ljubljana, pittsburgh, and tokyo: An experimental study. *The American Economic Review*, 81(5):1068–1095, 1991.

M. Rozendaal. Objects with intent: A new paradigm for interaction design. *interactions*, 23(3):62–65, April 2016.

E. Rukzio, J. Hamard, C. Noda, and A. De Luca. Visualization of uncertainty in context aware mobile applications. In *Proceedings of the 8th Conference on Human-computer Interaction with Mobile Devices and Services*, MobileHCI '06, pages 247–250, New York, NY, USA, 2006. ACM.

A. G. Sanfey, J. K. Rilling, J. A. Aronson, L. E. Nystrom, and J. D. Cohen. The neural basis of economic decision-making in the ultimatum game. *Science*, 300(5626): 1755–1758, 2003.

P. Scerri, D. V. Pynadath, and M. Tambe. Towards adjustable autonomy for the real world. *J. Artif. Int. Res.*, 17(1):171–228, September 2002.

G. Seyfang. Ecological citizenship and sustainable consumption: Examining local organic food networks. *Journal of rural studies*, 22(4):383–395, 2006.

T. B. Sheridan and W. L. Verplank. Human and computer control of undersea teleoperators. Technical report, DTIC Document, 1978.

P. Smith, C. McCoy, and C. Layton. Brittleness in the design of cooperative problem-solving systems: the effects on user performance. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 27(3):360–371, May 1997.

M. Soegaard and R. F. Dam. *Encyclopedia of Human-Computer Interaction*. Interaction Design Foundation, 2013.

T. Sohn, K. A. Li, W. G. Griswold, and J. D. Hollan. A diary study of mobile information needs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 433–442, New York, NY, USA, 2008. ACM.

T. Sohn, K. A. Li, G. Lee, I. Smith, J. Scott, and W. G. Griswold. *Place-Its: A Study of Location-Based Reminders on Mobile Phones*, pages 232–250. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

C. Southern, Y. Cheng, C. Zhang, and G. D. Abowd. Understanding the cost of driving trips. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 430–434, New York, NY, USA, 2017. ACM.

J.-Y. Sung, L. Guo, R. Grinter, and H. Christensen. "my roomba is rambo": Intimate home appliances. In J. Krumm, G. Abowd, A. Seneviratne, and T. Strang, editors, *UbiComp 2007: Ubiquitous Computing*, volume 4717 of *Lecture Notes in Computer Science*, pages 145–162. Springer Berlin Heidelberg, 2007.

S. C. Sutherland, C. Harteveld, and M. E. Young. The role of environmental predictability and costs in relying on automation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 2535–2544. ACM, 2015.

J. Tullio, A. K. Dey, J. Chalecki, and J. Fogarty. How it works: A field study of non-technical users interacting with an intelligent system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 31–40, New York, NY, USA, 2007. ACM.

A. van Wissen, Y. Gal, B. Kamphorst, and M. Dignum. Humanagent teamwork in dynamic environments. *Computers in Human Behavior*, 28(1):23 – 33, 2012.

J. K. M. Verame, E. Costanza, J. Fischer, A. Crabtree, S. D. Ramchurn, T. Rodden, and N. R. Jennings. Learning from the veg box: Designing unpredictability in agency delegation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 447:1–447:13, New York, NY, USA, 2018. ACM.

J. K. M. Verame, E. Costanza, and S. D. Ramchurn. The effect of displaying system confidence information on the usage of autonomous systems for non-specialist applications: A lab study. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 4908–4920, New York, NY, USA, 2016. ACM.

J. Wang, M. Lewis, and P. Scerri. Cooperating robots for search and rescue. In *Agent Technology for Disaster Management Workshop at AAMAS*, volume 6, 2006.

M. Weiser. The computer for the 21st century. *Scientific american*, 265(3):94–104, 1991.

L. Yang, Y. Cui, F. Zhang, J. P. Pollak, S. Belongie, and D. Estrin. Plateclick: Bootstrapping food preferences through an adaptive visual interface. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, CIKM '15, pages 183–192, New York, NY, USA, 2015a. ACM.

R. Yang and M. W. Newman. Learning from a learning thermostat: Lessons for intelligent systems for the home. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '13, pages 93–102, New York, NY, USA, 2013. ACM.

R. Yang, M. W. Newman, and J. Forlizzi. Making sustainability sustainable: Challenges in the design of eco-interaction technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 823–832, New York, NY, USA, 2014. ACM.

R. Yang, D. Pisharoty, S. Montazeri, K. Whitehouse, and M. W. Newman. How does eco-coaching help to save energy? assessing a recommendation system for energy-efficient thermostat scheduling. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '16, pages 1176–1187, New York, NY, USA, 2016. ACM.

R. Yang, E. Shin, M. W. Newman, and M. S. Ackerman. When fitness trackers don't'fit': end-user difficulties in the assessment of personal tracking device accuracy. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 623–634. ACM, 2015b.