



DELIVERABLE 6.6

A tentative proposal for an ethical framework

Author(s):	Heather Draper and Tom Sorell
Project no:	287624
Project acronym:	ACCOMPANY
Project title:	Acceptable robotiCs COMPanions for AgeiNg Years

Doc. Status:	Final draft
Doc. Nature:	Deliverable
Version:	2
Actual date of delivery:	8th October 2014
Contractual date of delivery:	Month 36
Project start date:	01/10/2011
Project duration:	36 months
Approver:	Zuyd

Project Acronym: ACCOMPANY

Project Title: **Acceptable robotiCs COMPanions for AgeiNg Years**

EUROPEAN COMMISSION, FP7-ICT-2011-07, 7th FRAMEWORK PROGRAMME
ICT Call 7 - Objective 5.4 for Ageing & Wellbeing

Grant Agreement Number: 287624



DOCUMENT HISTORY

Version	Date	Status	Changes	Author(s)
1	23/09/2014	Draft	N/A	Heather Draper & Tom Sorell
2	08/10/2014	Submitted version		Heather Draper & Tom Sorell

AUTHORS & CONTRIBUTORS

Partner Acronym	Partner Full Name	Person
UB	University of Birmingham	Heather Draper
UW	University of Warwick	Tom Sorell

Contributors:

Sandra Bedaf, Zuyd, reviewed the document

Short description

As required by Task T6.8, this deliverable evaluates the ethical framework proposed in D6.2 (resulting from work in T6.4) in the light of the results of the work with potential users in T6.7 (reported in D6.4). It first discusses each of the six values in the light of the data and then turns to participants' responses to the tensions suggested by the scenarios. It concludes that the value framework does not so much need revision as operationalization, and suggests some possible content for an agreement between care-robot users and those who introduce care-robots into older people's homes.

Table of Contents

Short description	2
Table of Contents	3
1 Introduction.....	5
2 The six values discussed in the light of the results of the ACCOMPANY user panels.....	7
2.1 Autonomy	7
2.2 Independence	11
2.3 Enablement	12
2.4 Safety	15
2.5 Privacy	17
2.6 Social connectedness	20
3 Participants' responses to the tensions.....	23
4 Operationalising the framework through prior agreement	26
4.1 Does the framework require revision?	26
4.2 Operationalizing the framework by an agreement	26
5 Conclusions	28
References	31

D6.6 A TENTATIVE PROPOSAL FOR AN ETHICAL FRAMEWORK

1 Introduction

In D6.2 (*Identification and discussion of relevant ethical norms for the development and use of robots to support the elderly in their own homes*) we discussed values that might be used to guide the development of care-robots for cognitively unimpaired older people who had hitherto lived independently in their own homes, but who, without additional support, would be unable to continue to do so. Six values were suggested.

- autonomy – being able to set goals in life and choose means;
- independence – being able to implement one’s goals without the permission, assistance or material resources of others;
- enablement – having or having access to means of realizing goals and choices;
- safety – being able readily to avoid pain or harm;
- privacy – being able to pursue and realize one’s goals and implement one’s choices unobserved;
- social connectedness – having regular contact with friends and loved ones and safe access to strangers one can choose to meet.

It was noted that circumstances could readily be identified in which two or more of these values were in tension. Given the target user group, however, and the relevant public policy aim of enabling older people to run their lives in the same way as other adults for as long as possible, we proposed that the overriding and organising value should be autonomy.

In D6.4 (*Report on data analysis aspect of ethical evaluation*) we reported the results of ACCOMPANY user evaluation panels who were asked to discuss realistic scenarios reflecting potential tensions between some of the six values. We wanted to see whether potential users – older people and those caring for older people formally or informally – would appeal to similar values when considering the scenarios, and how they might be inclined to resolve the tensions between them. These scenarios also left room for users to identify other relevant values. The scenarios can be found in table 1 below.

Table 1 – scenarios used in focus groups

Scenario 1

Marie, who is 78 years old, has lived alone since her husband died ten years ago. She has ulcers on her leg, the dressings for which are changed by a nurse once a week. It is important for the healing of these ulcers that she moves around as much as possible to encourage circulation to her legs and avoid further swelling. Her Care-O-bot® knows that she should be encouraged to move about, and suggests several times a day that she walks with it to look out of the window at either the garden or the street below. Marie is reluctant to get up from her chair because she is afraid of falling and walking is uncomfortable. She also uses the Care-O-bot® to get drinks for her from the kitchen, even though the nurse has suggested that she should go to the kitchen with the Care-O-bot® but let it carry the drinks back to her chair for her. Also the Care-O-bot® can only bring bottles of water to her and the nurse suggests that she would feel warmer if she made herself hot drinks. The Care-O-bot® reminds her to take her antibiotics and to keep her leg up on a stool when she returns to her chair after, for example, going to the toilet. She is grateful for the reminders about the antibiotics but feels irritated about the reminders to elevate her leg as she hardly ever forgets to do this but she likes to get comfortable first. She sometimes put her leg down so that her cat can sit on her lap more comfortably. Her ulcers are slow to heal but when the nurse asks if Marie

is moving around more she always says that she is, even though she ignores the prompts to come to the window and doesn't go to the kitchen with the robot.

Scenario 2

Frank is 89 years old and generally frail. He lives alone and needs assistance from a Care-O-bot® to live independently. He prefers the Care-O-bot® to having the neighbours or carers helping him because he thinks they are inclined to be intrusive and interfering. He uses his Care-O-bot® interface to talk about fishing with a friend he has known since childhood. Neither of them can go fishing anymore, but they enjoy talking about when they did and discussing items in a fishing magazine that they both subscribe to. They talk about once a month. Frank really looks forward to these conversations and they put him in a good mood for days afterwards. He becomes quite miserable if his friend is in hospital and unable to talk to him. Frank's daughter has suggested that the Care-O-bot® should be used to encourage Frank join a virtual fishing forum on the internet. She is worried that he only has one friend who is older and poorly and may die leaving Frank with no one else to talk to about fishing. Frank says that he is too old to be making new friends.

Scenario 3

Nina who is 70 years old had a stroke two years ago but has now recovered the use of her arm though one side of her face droops slightly. She is self-conscious about this, but it does not affect her physical functioning. She is supported at home by a Care-O-bot®. Since having the stroke she has become quite irritable and impatient. She often shouts at her daughter when she visits and complains angrily about her condition. Her daughter finds this very upsetting and has come to dread her visits. Nina has been so rude and demanding that two cleaners have already refused to work for her anymore. She is usually polite with her friends. Her Care-O-bot® has been programmed so that it will not do things for her if she asks sharply or in a demanding tone. It encourages her to say please and thank you and will withdraw help until she does so. Nina finds this infuriating and insists that the Care-O-bot® is reprogrammed to do what she asks no matter how she asks for help.

Scenario 4

Louis, who is 75 years old, is determined to continue to live in his own home, which is in a small town in which two of his sons live. He is regularly visited by his daughters-in-law, who bring him food, help with his cleaning and do his laundry. Louis was left with some weakness in one of his legs as a result of an accident in his 40's. He is becoming frail and is finding it increasingly difficult to get up from his chair and walk with his sticks. Louis is supported at home by a Care-O-bot®. The Care-O-bot® is programmed to help support him when he gets up from his chair and can be summoned to help if he falls. Louis has discovered that he can use the interface on the Care-O-bot® to visit online gambling sites and enjoys playing poker in the evening. He also uses the interface to give his doctor his blood pressure measurements, and sometimes his medication is adjusted as a result of the measurements he gives. Louis falls over about once a week on average. On the whole he is able to get up again with the help of the Care-O-bot®, but he recently was on the floor for several hours unable to get up and developed a bladder infection from lying in the cold unable to reach the toilet. He was in bed for several days as a result. This placed an additional burden on his daughters-in-law, who took turns to stay with him during the day until he was well enough to live alone. It was during this time that his daughters-in-law realised that he used the Care-O-bot® to play poker on line. They are very unhappy about this as he often loses money. They want access to the poker site to be blocked. They have taken away his sticks so that he has to use his walking frame, which means that he is less likely to fall. They want the Care-O-bot® to be programmed so that it alerts them as soon as he falls. Louis insists that it is up to him what he does with his own money and says that he doesn't want them to come rushing around every time he falls because he can usually get himself up.

Working with other partners involved in WP6, we have published some of the findings that emerged from our interpretation of the rich data gathered from the user panels (Draper et al 2014a, Draper et al 2014b, Draper and Sorell 2014, Jenkins & Draper 2014, Bedaf et al (in progress)). In this deliverable we shall discuss findings reported in D6.4 only insofar as these affect the question of whether the value framework needs to be revised. We shall first discuss each of the six values in the light of the data and then turn to participants' responses to the tensions suggested by the scenarios. We will conclude that the value framework does not so much need revision as operationalization, and we suggest some possible content for an agreement between care-robot users and those who introduce care-robots into older people's homes.

2 The six values discussed in the light of the results of the ACCOMPANY user panels

2.1 Autonomy

Autonomy in the sense of the value framework outlined in Deliverable 6.2 is the capacity to make choices and lead one's life as a normal adult. Autonomous adults are permitted to make irrational and even immoral choices within limits set by law. Older adults should not be assumed to be less autonomous than adults in general, and it would be seriously unjust to make choices for them or fail to consult them while not interfering with, and taking care to consult, younger people. Having one's own home both assists and requires autonomy, and care-robots are assumed, for the purposes of the value framework, to help people stay in their own home for as long as possible.

Participants in all three kinds of focus groups who responded to the scenarios thought that it was important to respect something amounting to the autonomy of older people in our preferred sense. Participants in the older people's groups tended to agree with our own view that age alone should not determine whether someone is free to choose how to live. Participants in all three types of groups drew a distinction between complying with the wishes of older people who had mental capacity (adequate cognitive function to make decisions)¹ and complying with the wishes of those who did not. The greater the mental capacity of the user, the more their wishes were to be respected; the less they had mental capacity, the less wishes were to be respected – especially where these wishes were thought by carers to be irrational or harmful, or where they depended for being fulfilled on the co-operation of carers. The same pattern of thinking was prompted where the carer role was played by the robot. So as long as the older person has mental capacity, robot refusal to accede to his/her wishes was generally opposed, as was reprogramming of the robot to frustrate the older person's choices. We are inclined to conclude that the participants were drawing a distinction between simply giving older people what they want irrespective of dependence and mental capacity, and respecting autonomous choices only once certain thresholds of mental capacity and independence were reached.

¹ We use mental capacity here as defined in UK law in the Mental Capacity Act 2005

But respect for the autonomy of the older householder with mental capacity was not always thought to be overriding, as we shall see when we come to tensions between autonomy and safety in particular. Participants also recognised that ethical problems might arise where the older person's exercising autonomy potentially erodes the autonomy of carers. This is evident from the way in which the participants in all groups referred to persuasion, negotiation and compromise in the care triad (older person-informal carer-formal carer).

Participants in all groups drew a distinction between persuasion and coercion. Here we use 'persuasion' to mean bringing about a willing change of attitude and 'coercion' to mean forcing someone to perform an action or accept some state of affairs against their wishes. Coercion is generally thought to be morally wrong – even when it brings about some benefit for the person coerced – because it undermines autonomy. People who never had or who have lost the mental capacity required for autonomous action cannot be wronged in the same way.² Persuasion, as we see it, engages with reason to effect change, though people might be persuaded by observation as well as reasoned argument; for example, they might attend a demonstration of technology capable of simplifying daily tasks like shopping, or they might see a friend using the technology and infer that they too were capable of using it after all. The line between persuasion and coercion can, however, be a fine one (Powers 2007).

Coercion might include forms of manipulation that are sometimes associated with so-called 'persuasive technology' in which subliminal psychological associations are sometimes harnessed to get an individual to change his/her behaviour in ways s/he wants but finds difficult, or in socially useful ways independently of his/her wishes (Berdichewsky & Neuenschwander 1999, Spahn 2012). We think that the ethical issues raised by these kinds of technology are not the same as those raised by a care-robot designed to encourage enablement or independence. One significant difference is that we are assuming that the care-robot is introduced into the householder's home with his/her agreement, and with the householder understanding its functions and role. This is not necessarily the case with e.g. manufacturer-installed programmes in cars to reduce fuel consumption or to control speed.

Participants in all groups regarded persuasion as an essential tool or skill in caring for older people, whereas paternalistic coercion was treated with more caution (depending what else was in play). Participants tended to be sceptical about how persuasive a non-human could be. They thought users might simply turn the robot off if they did not like what it was saying or perceived it to be nagging them. Again, dislike of a robotic, artificial voice and the robot's limited capacity to 'reason' with the older householder were regarded as important. On the other hand, participants also recognised the potential advantages of the robot in changing behaviour, namely that the robot can be programmed to persist and, as a machine, can persist without loss of patience or endurance.

² This does not mean that all forms of force in relation to non-autonomous people are therefore acceptable. The use of force can be harmful in other ways. Interfering in instinctive heteronomous desires can cause distress even though it does not undermine autonomy, for instance, as can trying to break patterns of behaviour that might be habitual without having been chosen. Force itself can also cause physical discomfort; indeed this physical discomfort may be the means by which the coercion is achieved since people will generally avoid painful stimuli. Using force on those who lack capacity can make them fearful and distrustful of the services that should be helping them.

When a robot persists it may not, however, be doing more than ordinary humans can do in the way of persuasion, it may be doing something else altogether. Robot persistence is more coercive than persuasive, when it does not give the user reasons for co-operating and even where reasons are given, wearing the older person down – rather than changing their mind – might be regarded as coercive. Nor is the robot able to engage in anything more than a simulation of discussion. For this reason the scope for genuine robot persuasion of the older user may be very slight or non-existent.³

Human persuasion can be paternalistic: it can aim at the good of the subject of the persuasion whether or not the subject cares about that good. For example, one person can try to persuade another to eat or drink less for the sake of their own health, when the subject cares much less about health than the satisfactions of eating or drinking. Coercion motivated by paternalism and enacted through the robot was generally, but not universally or unconditionally, disapproved of by participants in all groups. Although human persuasion, even fairly persistent persuasion, aimed at doing the user good was thought to be acceptable by many focus group participants, there was general agreement that the older householder could resist or reject that persuasion. This suggests that for our participants, adult autonomy can override outsider concerns for adult welfare, contrary to paternalism.

Participants in all groups seemed to regard subterfuge as a legitimate tool of persuasion, whereas in moral philosophy intentional deception is usually thought to undermine autonomy, and therefore to depart from persuasion in our narrow sense. Getting people to *try* new things in order to make an informed choice was often regarded as a reason to apply pressure or other coercive practices to older people; but only for a limited period of time. This response was particularly pronounced in relation to attempts to get Frank (Scenario 2) to increase his social contacts through the medium of virtual social networking. Autonomy-promoting paternalism (forceful intervention for the sake of getting someone to make their own choices) might therefore be acceptable, though participants tended to think that this kind of interaction would inevitably be human-human rather than robot-human.

As we shall see again in relation to tensions between autonomy and other values, many participants agreed with our view (D6.4, p. 18) that holding people to agreements willingly entered into (other things being equal) does not undermine their autonomy. For instance, some participants thought that it was not unreasonable to expect users to cooperate with the robot if they had agreed to have it in their homes. In a similar way, many thought that it was reasonable for the robot to be removed (for use by someone else who needed it) if it was not being used as agreed. All of this suggests that respect for autonomy is not just something that needs to be programmed into the robot, but into the policy for placing a robot in an older person's home.⁴ A policy that involves a well-understood agreement by an

³ A related note of caution: It is possible to envisage, as writers of science fiction have for many decades, the possibility of robotic 'persons' (in the philosophical sense of the term) whose autonomy should be respected. Our framework, however, is geared to a notion of a robot that is autonomous in the simplest sense of not operating directly under remote human control. We do not, therefore, assume that there is any reason to 'respect' the robot or its actions except insofar as the failure to do so impacts on other humans. Equally such a robot cannot be praised or blamed for its actions.

⁴ Lack of respect for autonomy in these contexts can result in adaptive preferences (Elster 1983, Walker 1995, Nussbaum 2000) that should not be confused with autonomous wishes. The willing acceptance into one's home

older person to a care regime involving a robot is at least autonomy-respecting. It can even be autonomy promoting if, as a result of the care regime, the user is able to make more choices.

The notion that users might be held to agreements entered into at the time a robot was introduced into their home was more prominent in the results than in our original discussion. It will therefore be discussed further in section 3 below.

All groups acknowledged that in the care triad (older person-informal carers-formal carers) the wishes and interests of people other than the older householder needed to be taken into account. Consideration of conflict of interests was prompted in particular by the fourth scenario, where the participants had been explicitly informed that Louis' reluctance to programme the robot to alert carers to his falling had resulted in him spending a prolonged period of time on the floor, which had in turn resulted in an increase in care needs that had to be met by his daughters-in-law. Of the three groups, the formal carer participants were generally less willing to settle conflicts in favour of the older person, though they were not especially sympathetic to the interests of informal carers. Rather they drew attention to the fact that they were themselves a limited resource that had to be distributed fairly among their clients.

The ACCOMPANY robot has tended to be conceived as a resource for a single householder; so participants were not prompted to consider how the attentions of the robot should be divided between multiple users, and this may have made it easier for the participants to reject more assertive actions by the robot, such as resisting a command. As Jenkins and Draper (2014) note, it is likely that some of the existing difficulties in the care triad will be played out through the robot and it is hard to see how design per se can be used to address some of these power struggles without designers taking sides.⁵ Given the target group of ACCOMPANY, however, there are reasons to privilege the autonomous choices of the older householder in conflicts, provided that the householder has had an effective say in how and why the robot has been introduced (see section 3 below), and provided that the householder is willing to accept any adverse consequences of their choices. In general – particularly when it comes to healthcare – it is not clear that individuals are in fact required to suffer the bad consequences of their choices. On the contrary, the widely accepted paternalism of medicine requires the same response to an emergency of the patient's own making as to an emergency that befalls the patient through no fault of his/her own. We have argued elsewhere (Draper and Sorell 2002) that this tendency to feel obliged to 'rescue' patients from the adverse and often foreseen effects of their autonomous choices may in fact count as an unacceptable paternalism, because it makes captives of doctors and other health care workers. We also argue that patients may be obliged to cooperate with healthcare workers (Draper and Sorell (2002)) when they have willingly sought their advice. A parallel conclusion is that older people

of a robot with specific features and for specific purposes distinguishes this sort of robot from other forms a persuasive technology discussed in the literature cited above.

⁵ The question of which side *should* be taken is an issue that is much wider than that of robot ethics, and is one that needs to be addressed independently given the extent to which older people are likely to be dependent on – and are likely to want to draw on – human care (formal and informal) for some time to come.

have to abide by the autonomy-promoting demands of a public authority's care agreement when the care agreement introduces a robot into the household.

The participants were aware that the resources of the robot were not as limited as the resources of an individual carer; for instance, the robot would not get tired or worn down by unreasonable demands or unremitting care needs, and this was generally regarded as an advantage in terms of the provision of care. The scenarios did little to prompt the participants to consider directly that a care-robot has needs, and that a failure by a user to accommodate these could impact on the care the robot could provide. For instance, the participants did not take into account the fact that without experiencing fatigue the robot would need literally to recharge its batteries in order to continue to function, and that it is likely to be programmed to return to its charging base when power levels are critically low rather than grinding to a halt mid-task. In hindsight, it may have been useful for this type of resource limitation to have been factored into at least one of the scenarios.⁶

2.2 Independence

According to the value framework suggested in D6.2, people are independent when they are able to act on their choices without significant help from others. The ACCOMPANY care-robot is not for the incapacitated or seriously disabled but for those who as a result of increasing age-related frailty find it harder to carry out certain tasks, e.g. lifting or house-cleaning. Independent older people might be able to carry out these tasks while taking longer to do them than younger people, perhaps much longer than they would otherwise choose to spend on them. On the other hand, they do not depend on others to decide on their activities, or to feed and clean themselves, or to take medication.

Views about independence (as distinct to autonomy) were not especially prominent in the focus group results. A few participants noted that the way in which a householder chose to use the robot could ultimately erode its ability to promote independence. They noted, for instance, that the fetching and carrying functions may disincline users to fetch and carry for themselves, with the possible result that they lose the ability to fetch and carry for themselves and so require more care. The potential for a care-robot to widen the care gap is explored further by Bedaf et al (in progress). It reflects a tension between independence and autonomy. A care-widening care-robot may well aid autonomy by facilitating the realization of older person's choices, but it can at the same time deskill him/her.

Whether independence gives way to autonomy may depend on how the robot is introduced and how its use is agreed with the user. Older people might autonomously accept a robot on the understanding that it will help them to do things themselves rather than have things done for them. If so, then their

⁶ This may, however, not have greatly affected the value placed on autonomy in the ethical framework, since in ethics it is accepted that the exercise autonomy may be constrained by external factors that are beyond the will or capacity of other moral actors to change. Potential care gaps left by a recharging (or malfunctioning) robot need to be considered, particularly urgent or unanticipated needs. Beyond this, it is not unreasonable, nor an erosion of autonomy, to expect users work around necessary constraints, much as they would need to accommodate the limitations of any other mechanical device.

autonomy is not undermined when the robot ‘refuses’ to do things solo⁷ for the older person, even if at times s/he finds that more convenient. The autonomously agreed contract trumps the lazy autonomous choice of the moment. This is because the extent to which individuals should be free to erode their own independence may be limited, especially in welfare states where dependency above a certain threshold generates needs that co-citizens have to meet through taxation revenue. It may be unjust to add to their burden through willing loss of independence. It may even be unjust to expect family members to meet the needs generated by willing loss of independence. In any case, the ACCOMPANY robot is conceived as a means to reablement as well as an aid to implementing autonomous choices. So it may not be designed for a care agreement that allows for the easy surrender of independence.

The participants seemed to favour a balance between independence and autonomy. For example, they generally supported the idea of a care-robot designed to give reminders to take medication. Difficulties remembering to take medications – due to degrees of memory loss or complexity of medication regimes – are an impairment to living independently, and therefore having reminders was regarded as useful support, but something that fell short of the take-over by other people of the administration of medication. Reminders associated with physical activities associated with rehabilitation were understood differently, and these will be discussed in the next section.

2.3 Enablement

Enablement in the sense of D6.2 is a process, possibly involving the care-robot, of acquiring or regaining a skill needed for carrying out daily tasks oneself.

Attitudes to enablement were mixed. Participants could see the value of a robot that was able to offer rehabilitation and also make new skills and activities available for older people. Views were expressed that can be interpreted as concerns about the scope of enablement and the potential conflict with autonomy. Our data did little to settle the question of whether enablement was valuable in its own right, as opposed to being a potentially useful property of the robot.

Concerns about enablement tended to be bound up with the robot’s capacities for persuasion as opposed coercion. As already mentioned, participants expressed doubts about the ACCOMPANY robot’s ability to persuade, and they were concerned about the robot forcing cooperation from its user.

The scenarios provided different examples of enablement that the ACCOMPANY robot might support. In the first scenario Marie was reminded by the robot to take medication (see previous section), to keep her legs elevated and encouraged to move around. In the second, Frank was being encouraged to widen his social connections (responses here will be discussed in section 2.6). In scenario three, Nina was being encouraged to be less rude to her carers. In the final scenario, Louis was using the robot as a portal to a gambling website and to help him to get up when he fell. Although we had envisaged that discussions around enablement would be prompted mainly by the first scenario,

⁷ Doing things with the robot helping is one way of in which the ACCOMPANY tries to promote reablement for older people in its design.

the participants made comments of interest in response to all of the scenarios, which we shall take together in our discussion.

The participants seemed to envisage a spectrum of health-related enablement, with reminders to take prescribed medicine at one end and health-promotion at the other. In general, enablement was regarded as a form of behaviour modification, and enablement interventions were placed on the spectrum according to whether participants thought that a particular behaviour should or should not be modified. So, reminders to take prescription medicine was regarded a relatively uncontroversial, whereas using the robot to prevent smoking, alcohol consumption, physical inactivity and poor diet were more controversial. They were more controversial because they involved interfering with people's liberty to take risks with their own health, in particular through the intervention of a robot that was supposed to help, not control, its user. Comments made by participants suggested a particular concern over the possible continual presence of the robot, and the possible continual monitoring of the user with a view to interference. Some participants seemed to worry that permitting the robot to modify the behaviour of users at all would be the start of a slippery slope leading to the robot's taking control. Other participants were concerned about the robot's interfering in possibly harmful but nevertheless autonomous choices expressive of the user's strong or characteristic preferences. Participants might not have approved of the choices individuals made in the scenarios (Frank's gambling was considered reckless by some, for instance) but they regarded preventing certain choices as an attempt to change what someone was like. This they generally disapproved of, particularly in relation to the third (Nina) scenario, as we shall see. For many different reasons, then, the participants often seemed to favour autonomy over interventions for the sake of enablement.

In fact the tension between autonomy and enablement may be more complicated. Participants did not disapprove of efforts to enable older people; what they were concerned by is how far these efforts would be made by the robot as opposed to human carers who were able to negotiate with the older person. In each of the groups, negotiation and/or persuasion was regarded as completely acceptable, so long as rejection of suggested behaviours was open to the older person.⁸ Participants were concerned about whether the robot would be so inflexible as to be coercive. As mentioned above in relation to autonomy, the groups tended to doubt that a robot could replace a human when it came to coaxing a person. The older people's groups in particular were concerned that in the case of Marie, the robot might not recognise and give sufficient weight to the fact that movement might be painful, or more painful on some occasions than others. They disliked the idea of the robot somehow forcing older people to perform painful actions, even where these were part of a therapeutic regime.

In papers disseminating the results of our work for this project, we have already alluded to the possibility of looking at this issue from the perspective of ethical norms for physiotherapy (see Sorell & Draper (2014), Draper & Sorell (2014)). The process of rehabilitation requires an effort of will on

⁸ Indeed, the some participants seem to favour weak paternalism (Childress 1982) in some circumstances. For instance, some thought that removing Louis' sticks and providing him with a frame was acceptable because Louis could still get around using the frame, or even chose to try to walk without it (it would just be more difficult for him to do so). Likewise, some thought that the robot could be reprogrammed so as not to be a portal to gambling sites because this still left open other avenues for gambling.

the part of the person seeking to be re-enabled, and this may include making an effort in the face of physical discomfort and frustration created by the inability to perform an action, particularly an action that could previously be performed with ease. Technology sometimes accommodates more passive rehabilitation (as in the case of mechanical devices that gently and repeatedly move limbs to rebuild muscle strength and increase movement range) but even these may require the user to make some effort and endure some discomfort. Such devices, although they are set up by physiotherapists, remain in the control of the user; if the movement causes too much pain, the user may simply stop using the device. The question is whether the older person could and should have a similar level of control over a care-robot. In our scenarios and topic guide, the scope for the robot to control the older person was very limited. It could verbally encourage movement ('come to the window') or perhaps resist a command (refuse to fetch something to encourage the person to get it (move) for her/himself). The robot was envisaged turning off a TV or positioning itself in front of a TV until the Marie elevated her leg. Even though the participants were not averse to robotic enablement, they disapproved of the robot's seemingly asserting itself. Sometimes participants appeared indirectly to express a fear that a robot might literally force someone to perform painful movements, which they regarded as dreadful.

Human enablers might also be fairly forceful in their attempts at rehabilitation; is their being forceful also dreadful? Human force (within the limits of enablement) may be justified by its net benefits (in keeping with what we elsewhere call the 'no gain without pain' principle Draper & Sorell (2014)). But the harmful effects must be minimised, and the patient must agree to both the ends and the means. It is not obvious that the moral character of the situation is changed by substituting a robot for a human enabler, provided similar safeguards are in place. This does not mean that enablement trumps autonomy, but rather that enablement, even painful enablement that is autonomously agreed to, is permissible. Provided that the older person could refuse to have an enabling robot in her home, and provided that she understood what the robot would be doing and why and still agreed to have one, it may not be wrong for the robot to 'take control' when the older person resists rehabilitation.

In any case, a user is always likely to have some powers of veto over the robot. This is partly because the backdrop for the development of social robotics actually privileges safety over all other considerations. Under the current regulatory systems of most Western countries risky or uncomfortable robot-human experiments conducted healthy volunteers would be treated as ethically suspect. Even more unthinkable are experiments on a population that is in some way vulnerable, and that entail the robot causing discomfort in the face of participant resistance, even with the participant's consent. We disagree with this strict safety regime, so long as discomfort is consented to. In D6.2 we argued that autonomy should be the organising value within the ethical framework. In our view, systems that do not let older people take some risks with safety may be regarded as unjustifiably paternalistic, especially in cases where the risks are consented to. There may be even greater reason to allow older people to take risks for the sake of their own enablement.

What should happen if a user agrees to have an enabling robot and then will not cooperate with its enabling functions? In such circumstances, it is not unreasonable for the funding authority to remove and reallocate the robot even if this means that the user thereby loses access to other functions the robot was providing. Many of our participants, as we have mentioned above, supported this position. We will discuss this further in section 3.

In D6.2 we considered whether some kinds of care-robots were infantilising. The literature we discussed was mainly preoccupied with robots that looked like children's toys. Our data suggested another potentially infantilising presentation of enablement to older people. This is what we describe as the 'let's do it together' method of coaxing older people to try new things or engage with enablement. This type of infantilisation does not consist merely of a tone of voice (the sing-song tone often used by adults to address a child). 'Let's do it together' coaxing is infantilising because it may fail to engage with the older person's reasons for not wanting to perform an action or behave in a particular kind of way. It may indeed assume an absence of reasons for not co-operating, questionably positing instead a kind of older person's stubbornness corresponding to childish refusals to co-operate. Adult-to-adult persuasion operates under a norm of giving reasons to a person which, if accepted, justify the choice of co-operating and make the co-operation autonomous. If one is seeking to persuade someone, one tries to identify and take seriously any reasons for disagreeing. One does not just assume that the obstacle is stubbornness or timidity born of having to try something new without support. Only if that sort of timidity *is* operating is the 'Let's do it together' strategy *not* infantilising. And arriving at the conclusion that the older person is timid ought to (morally ought to) proceed after an attempt to identify an articulate reasons for co-operation or non-co-operation.

There may be occasions where 'doing it together' is unobjectionable. If someone feels unable to walk in the park because they are afraid of tripping, then offering to walk with them and lend a supportive arm addresses their fear. It takes it seriously and offers a potential solution to a problem that is reducing the choices available to the older person. But if, on the other hand, someone says that they do not care for walking in the rain, offering to get wet with them misses the point. 'Let's do it together' may suggest that, like the child, all the older person requires to change their mind is an encouraging presence while they get on and do something they really do not want to do. When adults form supportive pairs or groups those banding together all want the same thing and feel that they are offering mutual, not patronising, support to achieving an end in which all share; they are doing together what they would struggle to achieve alone. It is a form of solidarity. 'Let's do it together', on the other hand, may be an offer the only aim of which is getting the other person to do something they do not want to do. It is often something the person doing the offering is already able to do effortlessly – it is not necessarily a declaration of solidarity. In this respect, robotic efforts of the 'let's do it together' variety might always be patronising since the robot is not capable of appreciating the end, whatever it is programmed to say by way of encouragement along the way.

2.4 Safety

Safety is being insulated from sources of harm. The insulation can be provided by one's own choices and policies or by the interventions and policies of others. There is a significant but difficult-to-locate line between minor and serious harm and this is particularly important for our purposes. Life-threatening harm is unlikely to be outweighed by any benefits; and so it can have priority even where autonomy is at stake. Incapacitating harm, especially foreseeable and avoidable incapacitating harm, also has a lot to be said against it, even when the decision to risk it or suffer it is autonomous. The reason is that it increases dependence and may generate reasons for restricting the scope of choice that would otherwise have been available to the agent. In Deliverable 6.2 considerations like these were

cited in support of the claim that safety from serious harm was nearly as high in the hierarchy of our six values as autonomy itself.

The safety of the older householder was raised in response to all of the scenarios. It was a concern for some participants even where the scenario was not designed to emphasise safety. Some participants were concerned that harm could befall Marie and Nina if the robot refused to act on their instructions. Participants were also concerned about the potential dangers of internet interactions in the Frank scenario, and the risks of gambling in the case of Louis. In the Louis scenario as well, many were very uncomfortable about Louis being able to prevent robot alerts to outsiders if he fell or remained on the floor for long periods.

In the Louis scenario the user himself was in control of the programming and elected not to programme the robot to summon help, a decision that was questioned by his daughters-in-law in the scenario. In all the focus groups, the predominant feeling was that the robot should summon help in the event of a fall regardless of the older person's wishes to the contrary. The most commonly proposed compromise was that the householder be given time to get up before the robot alerted external agents. One prominent reason for this compromise was that it reflected the care role of the robot. Participants found it incongruous that a robot could be present and not summon help. Some participants tended to anthropomorphise the robot; it was almost as if it was a human being standing idly by and doing nothing. For others the robot represented a safety net that should not be disabled. Our participants mostly supported the use of a default alert setting when it came to falls: the householder could choose within narrow limits how quickly an alert was issued but would not be able to override the default setting completely. They thought that the user could also be given a choice about whom to notify – this might not be the daughters-in-law in the case of Louis, for instance – but they seemed to suggest that it would be unacceptable for *no-one* to be alerted. There was no specific agreement amongst participants about the precise point the alert would be sounded regardless of the users' wishes. Instead participants spoke vaguely about the point at which the user would suffer harm if help was not forthcoming. We return to this point at the end, where we discuss some factors that need to be taken into account in operationalizing a robot that defers to the user's wishes.

We have commented elsewhere (Draper & Sorell 2013) on the ethical difficulties that may arise when telecare technology can detect falls and older users can disable this equipment. Falls undoubtedly create demands on health services and can lead to longer term difficulties and health problems for older people – even those who up to the point of falling were fairly independent. Where these demands are made on resources in welfare states, it may be reasonable to ask or even require citizens to minimise these demands. This may mean not using services frivolously, taking precautions against infection, or adhering to advice and treatment regimes. In the same way, people might be asked to minimise demands on informal carers. If someone is dependent on the good will of others for help, this provides them with a reason not turn for help unnecessarily. Arguably, the more dependent one is, the greater the need for cooperation which prevents greater dependence or emergency dependence on informal carers. Co-operating with a robot care regime may be a case in point, but unless the care-robot can provide everything provided by informal carers, the interests of informal carers should play some part in negotiations leading to its installation into the older person's home (see section 4 below). The participants tended to agree with that line of thought. However, the interests of informal carers that they believed were relevant were only those directly related to the care they provide.

Different considerations were considered relevant in different circumstances. Participants did not believe informal carers had the right to frustrate older people's life-style choices, even if they cost money and threatened carers' inheritances. The issue was posed clearly by the Louis scenario, in which an older person was involved in online gambling, with all its risks of increasing dependence. Getting into debt was generally viewed as socially irresponsible, which may provide a reason for restricting the gambling stakes (e.g. by imposing the 'affordability' ceiling). If the robot is the medium through which socially irresponsible behaviour is facilitated, then modifying the programming to prevent such behaviour may be acceptable. On the other hand, the 'good will' that should motivate the provision of informal care might not be compatible with limiting spending simply because it erodes inheritance.

In all of these cases it is important to bear in mind that the limitations considered were not confined to older people but could apply to others who are much younger. As suggested in our first deliverable, the test should be whether proposed restrictions on the older person's behaviour could be justified to people who were not old, a middle-aged autonomous adult, for example. And the arguments proposed above cut across generations when it comes to going into debt and over-using welfare services. The arguments apply to anyone who is autonomous but dependent - regardless of age.

2.5 Privacy

A person enjoys privacy when there is restricted access to information about them, including information that can be gained by observation. Not every kind of collection of personal information is necessarily an intrusion. A thermometer does not invade someone's privacy when it measures their temperature, for example. The care-robot does not necessarily invade privacy when it is given an older person's schedule for taking medication. The central cases of invasion of privacy concern the unwanted or unauthorized collection of personal information by agents who might distribute or act on the information for their own purposes, and possibly against the interests of the person that the information concerns. If a care-robot that collects personal information with the users' consent were hacked by a stranger or even by a nosy carer or family member of the care-robot user, then the hacker rather than the care-robot would be invading the user's privacy. In any case, the invasion of privacy can sometimes be justified by the prevention of serious harm or by the requirements of a general care regime that the user has agreed to.

Focus group participants in general agreed on the importance of privacy. Our formal carer group tended to discuss privacy in relation to formulas and routines embedded in their professional codes of conduct and good practices. Other groups tended to describe their views in terms of unwelcome intrusion or 'Big Brother' surveillance but at the same time, there was little resistance to, and some positive support for, information being accessed by health professionals for direct therapeutic purposes. Here the robot seems to have been regarded as an extension of the healthcare professional or even itself a therapeutic tool. Participants were, however, concerned about health information being accessed by or passed to family members/informal carers. In this respect they seemed to be applying the norms of medical confidentiality.

The formal care participants in some groups were concerned that the robot could be used to monitor the care they were providing. Some seemed concerned that this could interfere with that care but

without elaborating on how. It is unlikely that formal carers could appeal to a right to privacy to object to monitoring of their working practices, especially if they were informed in advance. Some formal carers complained that they did not have sufficient time with their clients, and in a different context they explained that they sometimes had to spend more time than expected with one client, which made them late for an appointment with another. Additional time pressures created by care-robot monitoring may generate hostility to the robot unless this information is also used to improve their working conditions. Employing the robot to “police” care may discourage poor care practices, with benefits to older people. But it can also intrude on the privacy of the older person. Striking a balance between monitoring for good practice and privacy may be difficult where care involves nudity or captures private conversations. Recording would almost certainly require the consent of the older person, except where there were suspicions of both poor care and intimidation. Whether all care/medical interactions should be video recorded – and kept as part of a patient’s medical record – is already being considered in some jurisdictions. Such recordings could provide a definitive account of an interaction in the event of legal challenge, disciplinary action or unforeseen outcome. Such a policy raises complex data protection and access issues. For instance, access to recordings might only be granted for audit purposes or where there were suspicions about misconduct.

In our view, using the robot to police care would not violate the privacy of the formal carers.⁹ All care compromises privacy to some degree. Ensuring that appropriate care is provided may necessitate careful record keeping to ensure a smooth hand-over of care between carers, and so that care can be audited and improved. Carers themselves see aspects of a person’s life that they would not otherwise witness. Providing good care may depend in part on remembering these details, but even if it did not carers could not be *required* to forget them. Humans cannot will themselves to forget. Carers may be required to recount their experiences to others, or they may be required not to disclose them. At other times disclosure may be selective or heavily edited. The robot that records all its interactions with a user is in some senses similar to a human carer with a memory and does not therefore raise any greater concerns for privacy than human care does. The privacy concerns are raised by access to information. In this respect a gossipy and judgemental carer may be more invasive than a care-robot.

Our participants were unclear as to the information they thought the robot would be recording and in what format. Potentially, a robot could make continuous video recordings similar to a CCTV camera. Whether this would be privacy-violating would depend on why and how the recordings were made, what is recorded, who can access these recordings and on whose authority, how secure the data-storage system is, and how long the data is stored. Having visual images of robot-human interactions may, for instance, be useful for enablement. For example, it has been suggested that the robot may be able to enhance the user’s recall by providing pictures of when s/he last ate or drank, took tablets, telephoned a family member etc. (Ho et al 2013) Some of the formal care participants thought it would be useful to access information stored by the robot.

⁹ As the ACCOMPANY robot is targeted at older people living alone, we do not here consider the privacy implications for others – including informal carers – who might be sharing a home with the older person for whom the robot is being provided.

Consider, however, the responses of two formal carers reported in D6.4:

They could look at the print out together, that wouldn't be quite as invasive as the robot saying: 'Actually she didn't do that when I told her three times and she didn't get up!' (UH FC PF)

They cannot cheat, right?... That is the difference. The measures are taken and the robot sends them on to the physician. So there is no possibility to add a few degrees, or make it some degrees less. (ZUYD FC1 P2)

Looking at the print out together (with consent) could be a useful way of discussing and resolving obstacles to adherence to a care regime; the second quotation points to a different and perhaps questionable use of the monitoring record, namely to check the veracity of the patient's account. Where the robot collects data to enable a willing householder to be more independent, the data collection does not violate privacy. And more data may be therapeutically better than less. Consider, for instance, a robot that monitors whether medication has been taken and issues a reminder when it is not taken within a pre-specified period of time after it was scheduled. Such a robot might be more enabling than a robot that acts like an alarm clock and simply reports that now is the specified time to take the medication. In the former case, the user has the opportunity to remember for herself to take her medication, in the latter she may come to rely on the alarm rather than her own memory. She may be helped to live independently, but she can become dependent on the robot to provide the reminder as her capacity to use her own memory is eroded. To act as an alarm clock the robot does not need to collect personal information. To issue the reminder the robot needs to monitor what the user is doing. More information is stored (not just what medication should be taken and when, but also whether it has been taken), but with an enabling purpose. Assuming this information is only accessible to the older person (in the form of the enabling reminder to which she has agreed) her privacy is not violated. In the latter case, however, there may be a concern that the person will take the medication twice – once on her own volition and then again when the robot issues the reminder. It may therefore be argued that the robot needs to monitor whether the medication is taken in order to cancel an unnecessary reminder. Now safety concerns are beginning to resurface that appear conflict with the protection of privacy. Another concern might be that the robot should, like a human carer, be able to monitor medication adherence so as to alert someone if non-adherence reaches a dangerous level. This would be consistent with the position outlined above with regard to falls; the older person may want a higher threshold for intervention than carers are comfortable with, but the default position would still be that where the threat to safety was significant, the robot would be programmed to raise an alarm, and this ultimate threshold could not be overridden by the user.

Should the robot be able to alert someone if medication is not taken? Would this violate privacy? It may, and it may deprive the user of the liberty not to comply with a care regime. These concerns may nevertheless be outweighed by considerations of harm. In any case, some loss of privacy is going to occur as soon as one is cared for – by a robot or human alike. What adds to the privacy violation, if there is one, is the unauthorized but perhaps justified transmission of information to a third party.

The presence of the robot, like the presence of a human carer, unavoidably takes away some privacy. On the other hand, it may be an avoidable violation of privacy for healthcare professionals to have access to information stored by the robot (for the purpose of routinely monitoring adherence and

honest reporting of adherence – the kind of use suggested in the second quotation). Patients can be inaccurate or dishonest in reporting their adherence to a care regime as well as their intake of alcohol and calories etc. (Buetow et al. 2009). One response to this is for practitioners to be sceptical about patient reporting and adjust their judgements accordingly. This scepticism is caricatured by fictional TV character Dr Gregory House, whose approach is encapsulated by statements such as:

- “I don’t ask why patients lie, I just assume they all do”;
- "It's a basic truth of the human condition that everybody lies. The only variable is about what";
- “...when you want to know the truth about someone that someone is probably the last person you should ask”.

It could be argued that patient dishonesty should not be encouraged and that therefore programming the robot so as to prevent a sceptical Dr House from interrogating its data is to collude with patient dishonesty. It might be argued that just because other patients (like our middle aged autonomous adult comparator of the previous section) are able to get away with being dishonest, does not mean that older patients with a robot should be able to, and relevant difference is not *age* but the presence of the robot. The middle aged comparator *with* a robot would be similarly constrained.

On the other hand, this may be to overlook an important difference between robotic and human carers and companions: robots are not moral agents. One of the reasons privacy is compromised when one takes a carer or companion (or even a servant) into one’s home is that this person can neither avoid exposure to personal information, nor avoid making sense of the information to which they are exposed. There is a shared understanding of the potential normative implications of a carer’s seeing an unexpected person sharing the householder’s bed, overhearing a phone call to the betting office or alcohol retailer, or reading a letter dictated to a solicitor changing a will. If confronted by a wife or child asking questions about these events, a human servant/carer/companion has to make a normative judgement about the relative importance of infidelity, gambling, alcohol use and disinheriting a child compared with some prior agreement to maintain confidentiality. For the robot there is no such tension. This could be regarded by some older people as a potential advantage of having robot as a carer. The robot is neither nosy nor takes any kind of prurient interest in the data it collects. It is merely data. It is not a life-story that the robot secretly or otherwise passes judgement on. In this respect although the robot may be ‘all seeing’ it may be less intrusive than a human carer who is present less often.

Giving the older person control over who can access their personal data from the robot is the best way of protecting their privacy, and also conforms to the norms for data protection. This means that healthcare professionals should not be able to check the veracity of a patient’s reported adherence without that patient’s consent, and patients with a robot, therefore, have the same scope for deception (or “cheat[ing]”) as patients who do not need robotic care.

2.6 Social connectedness

Social connectedness is valuable because it alleviates unwanted loneliness. Older people are more likely than middle-aged adults to lose their friends or spouses through death. Some forms of disability and incapacity due to old age can also make their surviving friendships among old people less

valuable. Connections with people, including strangers with common interests, can fill the gap left by old friends. However, the same loss of ability that can make the care-robot useful can also stand in the way of going out to make new friends. Those older people who are not put off by the internet have an answer to this pattern of problems. They can stay in touch with friends and family or make new contacts through social networks or Skype. The ACCOMPANY care-robot incorporates a tablet theoretically capable of linking people to the internet, and the care-robot can assist social connectedness by keeping a record of anniversaries, birthdays and the schedules of those with whom the older person would like to spend more time.

The importance of social connectedness was reflected in the discussions of the older people and informal carer groups, but was not prominent in the formal carers group, which tended to concentrate on the way their interactions with older people could not be simulated or reproduced by care-robots. Participants from other groups tended to agree that at least some human interaction was irreplaceable. The importance of social contact and being part of a community was considered valuable quite apart from receiving care from humans. Some of the groups did, however, discuss how social connectedness provided a care safety net for older people. For instance, being integrated into a community meant that neighbours and others noticed deviations from normal behaviour (not opening shutters or not being seen out and about) that could indicate an older person in difficulties.

The Frank scenario was designed to elicit reactions to a potential tension between autonomy and social connectedness. In response, some participants drew a distinction between loneliness and being alone, recognising that not all people who are socially isolated actually want or miss human company. Nevertheless, participants were generally in favour of coaxing older people at least to try to remain socially connected. This suggests that people should not settle for loneliness by the mechanism of adaptive preference. On the whole, then, the data we gathered supported the approach to the tension between autonomy and social connectedness that we adopted in D6.2.

Social contact is a means of forming reciprocal care relationships that older people sometimes value (Birmingham Policy Commission 2014: 34, Sharkey 2014). Care relationships that grow out of social capital are sometimes regarded as an alternative to, or at least a way of supplementing, the formal care funded by the welfare state (Birmingham Policy Commission 2014). Reciprocal care arrangements – broadly defined – may also work better than others at promoting dignity and respect (Birmingham Policy Commission 2014). The scope for the robot to enter into reciprocal relationships of this kind is limited. Other data collected for the ACCOMPANY user evaluations, particularly in France (see D6.7), suggest that there users might come to regard robots as a sort of pet, so that maintaining the robot provides the user with responsibilities analogous to looking after a dog. Although we argued in D6.2 that a care-o-robot might have advantages over telecare by having a ‘presence and alleviating loneliness’, we were clear that this ‘presence’ should not take the place of other forms of social connectedness. At the same time, it is possible that interactions with the robot which help to maintain its smooth functioning may become ‘purposive’ and therefore meaningful in the sense that they enable the user to be active in helping the robot to help them. Older people derive well-being from purposive activities, as opposed to those that merely ‘fill time’.

In D6.2, we took social connectedness to be valued by older people, that is, to improve their perceived quality of life. We envisaged that through its tablet a care-robot might help its user to create and

maintain new virtual relationships, and to stay in contact with existing family and friends. Our participants discussed both virtual social connectedness and maintaining relationships using video calling and social networking. The participants who spoke in these discussions all seemed to be familiar with this use of the internet, found interactions using Skype/internet fora useful, and many readily likened the type of use proposed to Frank in scenario 2 to their existing use of personal computers/tablets. Reactions to purely virtual relationships tended to be guarded. In the older people's groups particularly, many participants were not convinced that virtual relationships were a substitute for what they termed 'real' relationships.

Undoubtedly, older people who do not or cannot use the internet will face increasing social exclusion in the future. It might, therefore, be useful for a robot to encourage the use of the internet for purposes that connect older people to social institutions and services as well as its use to maintain and form new more personal relationships. In this respect, the participants' distinction between 'real' and virtual interactions have less and less application. The internet is increasingly completely replacing impersonal but 'real' (in the participants' sense of the word) interactions with virtual or automated ones. The more that the population in general purchases goods and services using the internet, the less it seems that this method can be regarded as poor substitute for other methods. At the same time, we know that over a million older people in the UK go more than a month at a time without seeing or speaking to family, friends or neighbours (i.e. having meaningful conversations); 600,000 only leave their home once a week or less; and, an estimated 5 million older people consider their television to be their main form of companionship (Age UK 2014). In the absence of meaningful interaction any direct human contact may be endowed with significance. This may act as a deterrent to embracing any interventions (such as the use of 'virtual' interaction or robot carers) that remove direct human contact, even at the cost of further social exclusion. Being 'troublesome', 'complaining' or 'deliberately needy' may be an older person's only means of accessing human interaction or attention. Maintaining this undesirable state of affairs fails to respect the dignity of older people, and may need to be addressed independently of robots, as the robot itself may be regarded as another barrier to meaningful human contact.

At the same time, virtual relationships or virtual contact could be re-valued upwards. One French participant – whose views are not representative of participants at large – was puzzled by the attitude of others in his group to the use of the internet. He said:

The word virtual is used, and is used when a screen is involved. When you're on the phone with someone, the word virtual is never used to describe it. [Others interject: But it's the same thing] Yes it is, so why is it that we don't use the word 'virtual' when telephones are involved but do when there's a screen, whereas with a screen we actually add something and can see the person we're talking to? I've been wondering about this for some time, I don't understand why. (MADoPA IF1 P1)

We have to agree: it is difficult to imagine an older person living alone without access to a telephone, or objections being made to installing a telephone on the ground that telephones discourage genuine social interaction or lead to social isolation and loss of meaningful relationships.

Undoubtedly, our participants were worried about policies of replacing human with robot contact. Harnessing robots to extend social connectedness will only go some way to addressing these concerns. The incentives are certainly present for designers to create affordable robots that can simulate certain human care functions. Who is ultimately responsible for ensuring that the robots they create are not, or cannot, be used to replace meaningful human contact is an open question.

3 Participants' responses to the tensions

As we have already suggested, participants generally tended to resolve the tensions presented in the scenarios by prioritising the autonomy of the householder. Where this was not the case, the tensions were supposed to be resolved by:

- (1) processes of negotiation, persuasion and compromise giving weight to all values and all personal interests in tension; or
- (2) assigning specific roles to the robot and then discerning an order of priority of values related to the requirements of fulfilling that role (e.g. carer, servant); or
- (3) by hypothesising an agreement under which the robot was introduced into the older person's home under certain conditions or for certain purposes that would give priority to one or another value in a given scenario;

(1) has already been discussed at length in section 2.1. Appeals to negotiation processes tended to privilege autonomy by assuming that ideally coercion could be avoided by reaching an agreed compromise. The participants accepted that not all tensions could be resolved by agreement; in which case, with some exceptions, they tended to resolve the tension in favour of the autonomy of the older person. In the case of the exceptions, participants tended to resolve the tensions with reference either to norms associated with roles attached to the robot or with reference to the agreement under which the robot was introduced into the older person's home.

As far as (2) is concerned, the scenarios were deliberately worded to *avoid* assigning a particular role to the robot. Most participants had already had some direct contact with the ACCOMPANY robot as part of the on-going user evaluations based at ZUYD, MADoPA and UH. Those recruited by UoB were shown a video of the robot. All had basic information about the ACCOMPANY project, the full title of which uses the word 'companion' in relation to the robot. This does not, however, appear particularly to have affected the perceived role of the robot in this round of evaluations. When the participants were unsure about how the robot should behave, or wanted to justify particular intuitions that ran counter to respecting the autonomy of the older users, they tended to refer to the norms associated with particular roles that they attributed to the robot. The most frequently applied roles were servant, healthcare provider, extension of a healthcare provider, as well as companion. Sometimes intuitions were based on the robot being merely a machine (just like any other household appliance), and at other times on its being a fairly sophisticated machine (more than just a common household appliance). No consistent role was applied but rather a role was applied that enabled the participants to argue towards a particular conclusion about which of competing values in a scenario should be given most weight.

We offer three examples to illustrate. First, when discussing the case of Louis, the participants were generally unable to countenance the idea that, at their request, older people should be left lying on the ground unable to get up when a robot was present that could summon help. The implicit reasoning here seemed to be that since no (human) carer would stand by passively after a fall, even at the request of the older person, the robot should not either, if its role was that of a carer. Similarly, in the case of Louis' gambling, participants wanted to protect Louis from the damaging effects of getting into debt and also to insist on his right to spend his money as he chose (even if that meant gambling it away). The tension was side-stepped by asserting that the robot *qua* healthcare professional should not facilitate or introduce Louis to gambling. If Louis found another way to gamble, then that was a different matter, but the robot was not an appropriate means. In the case of Nina, the robot was cast as a servant or machine that should not question user commands, even if the user was reprehensibly rude.

Using mechanism (3) participants addressed tensions in the scenarios by assuming that certain things would have been agreed by all affected by the robot's introduction into the older person's home. It was not that the participants came to the focus groups with a specific detailed agreement in mind, but rather that they thought the robot would be introduced under pre-agreed terms and conditions of some kind or other. We found this especially interesting because in D6.2 we suggested something similar:

Returning someone to a state of greater independence is certainly compatible with autonomy; the question is whether it is compatible with autonomy for a carebot to coerce someone to adhere to regimes that will return them to greater independence. The answer to this question may lie in what was agreed with the elderly person at the time a rehabilitation device or robot with enabling capabilities was provided. In the case of single-purpose device, there would be no objections to removing a state-funded device that was lying unused or not being used properly. Carebots pose a different challenge because they are designed to be multi-functioning and these other functions would also be lost if they were removed. Enablement functions are not quite the same as those providing potential social interaction. Disliking social interaction and preferring isolation is an issue of taste. Working against a carebot programmed to maintain independence is not simply an expression of taste, but a kind of resistance to independence. At the same time, the robot and its developers would not be working with the autonomy of the elderly users if the robot refused to do things that the elderly person could reasonably do for herself, or which it might be good for her to do for herself. (D6.2 p.18-19)

We developed this idea further in a paper reporting some of the results of the older people's focus groups:

It is significant that at least some participants volunteered the view that where robots were being provided for specific reasons (e.g. enablement) that the older person had agreed to, then the cooperation of the older person with enablement was part of the agreement. Although the participants generally favoured persuasion over coercion, some also felt that it was consistent with respect for autonomy to expect householders to deliver on their side of the deal. This view is consistent with that discussed elsewhere by Draper & Sorell suggesting that patients do have responsibilities, in particular, to follow through on care that has been voluntarily sought. What requires further consideration, however, is how providers of social robots with e.g. enablement features should enforce the householders' obligation to cooperate with those features.

If the norms of healthcare ethics are applied, householders would have to be provided with the opportunity to refuse continued consent to care. Anyone who really does not want to work with

the robot should not therefore be compelled by the robot or anyone else to cooperate. A failure of cooperation could, however, reasonably lead to the robot being withdrawn, regardless of how desirable to the householder other features of living with the robot might be (e.g. its ability to act as a servant to some degree, or even its providing some form of companionship, features that a simpler and less expensive robot might possess). Autonomy is compatible with having difficult decisions to make, and also with accepting the consequences of one's actions. On this basis an otherwise reluctant householder may feel compelled to live with what s/he regards as the less desirable features of the robot in order to retain the robot in his/her house so as to be able to keep the features that are valued. Provided that the enablement goals are not unreasonable and were understood in advance, such an outcome falls short of coercion and is compatible with respecting the autonomy of the householder.

Nor should cooperation be regarded as an adaptive preference under these circumstances. It is not that the user decides to settle for less autonomy by tolerating less control over information about e.g. falls or the absence of a veto on connecting with others socially; instead, it is an autonomous choice to accept the advantages of a companion robot alongside some policy-reflecting programming in the robot that goes against the individual grain. In the same way, one might autonomously accept a car, with its advantages, even if there were strings attached, like giving rides to the neighbours.

Accepting the policy-reflecting goals of the robot could also involve the autonomous acceptance of the withdrawal of the robot. One might envisage circumstances where a robot is withdrawn because its presence in the householder's home has achieved the enablement goals that were set for its installation. Loss of the robot in these circumstances is not unfair. It is on a par with returning crutches once a broken leg is mended. Nonetheless, the more generally useful and the more effective a companion the enabling robot is, the more likely it is to be missed by the householder when it is withdrawn. This is especially true of the withdrawal of multi-functioning social robots, which might be missed for their fetching and reminder functions, even their activity prompts are not. Dividing the same functions between different types of less complex robots might provide householders with a greater range of choices, and a more flexible and responsive service. (Draper et al 2014a)

We do not argue here that older people (or their formal or informal carers, for that matter) should be compelled to adhere to the agreements under which the robot is introduced into the older person's home. What we say is that it would not be unreasonable for a robot to be removed if it is not being used for its intended purpose, even if the removal disadvantages the user. In section 2.3 (above) we explored how forceful a robot could reasonably be if it is installed for the purpose of re-enablement.

If agreements with older people are to some extent to be regarded as binding, it is of course very important that this is made clear to older people in advance. Equally, we need to consider what alternatives might be available for those who cannot for various reasons agree to the terms for introducing the robot into a home. Agreements with older people are potentially empowering. For instance, in section 2.6 we supported the view that it is normally undesirable for care-robots to be used to reduce the amount of human-human contact older people have. If an older person discovers that introducing a robot is being used to justify a reduction in human contact that she did not foresee or agree to, then this would then provide her with grounds for rejecting the robot and insisting that her needs are met in ways that do not reduce meaningful human contact.

4 Operationalising the framework through prior agreement

In this section we first consider whether the results justify a revision of the value framework. We then enlarge on the way that an explicit agreement between user and robot provider could help to operationalize the value framework, and in particular its emphasis on autonomy.

4.1 Does the framework require revision?

The data we collected did not suggest that the six-value framework required significant addition or revision. Participants tended to recognise the importance of all of the values proposed without apparently calling attention to entirely new ones. The participants did refer to the financial cost and value for money of the robot. This may suggest that they thought that cost-effectiveness should be a further value, and this would be a consideration in any event where care-robots were publically funded. But not all robots will be publically funded. Given data that emerged in the final user evaluations in France, mechanisms that engage the older person with simple aspects of the maintenance of the robot may add an element of reciprocal care (care of the robot) to the care provided by the robot. In a different vein, the value framework could explicitly be interpreted as supporting autonomous decisions with ill effects on informal carers or friends or even state welfare services in cases where the effects of those decisions is to make users more dependent.

4.2 Operationalizing the framework by an agreement

Rather than a revision of the value framework, the results from the focus group seem to call for operationalization of its emphasis on autonomy. Operationalization includes, crucially, the processes by which the robot is introduced to the user's home in the first place. We have already referred to the focus group members' appeal to an agreement under which the robot is introduced. Deliverable 6.2 had independently pointed to the usefulness of such an agreement. The agreement would set out the purposes that e.g. a local housing, council or health authority had in offering a care-robot, and we are assuming that one of these purposes is the promotion of autonomous and independent living. If this is right, then the agreement needs to depend on processes in which potential users of the robot are (a) informed of the capacities of the robot; (b) consulted about which of these capacities might be useful to them; and, (c) informed of the options to refuse or withdraw co-operation with the robot in its exercise of capacities that the older person finds useful. The options to withdraw might themselves be activated after a trial period without those options, just to make the older person aware of what living with the robot might be like and how useful it could be. Similarly, there might need to be a trial process of withdrawal of the robot so that the older person can experience what life without the robot might be like if it is withdrawn. Ideally, potential users would be seen individually and face to face, with discussion encouraged. No less seems reasonable when so expensive a piece of equipment, and such an unfamiliar one, is introduced for long term use in someone's home. In addition to the older person's own interactions with the robot, the agreement would have to take into account data-retention by the robot and retransmission of the data to: (i) the robot-introducing authority; (ii) formal carers, including healthcare professionals; and, (iii) informal carers and family members. In keeping with the

value framework, the older person should normally be given the opportunity to veto data-sharing with certain groups listed or certain members of groups listed.

Beyond any trial period with the robot, a pattern of non-adherence could be allowed to develop up to a threshold where an interview about removing the robot was triggered. Allowing the older person a chance to see and discuss the evidence of non-adherence may be important to a subsequent decision to co-operate with the aims of enablement more whole-heartedly, or may prompt a reconsideration of the suitability of independent living, or may call attention to deficits in the original agreement that need to be resolved. Either way, the user has an autonomous choice to make.

The process so far outlined does not mention the possibility of a user's simply turning off the robot's monitoring functions, or the possibility of overriding an emergency alert. These possibilities are relevant to the older person's control over information about themselves, including information about falls. Control of such information is greater for unaccompanied older people than for accompanied ones, and loss of control might discourage some older people from opting for a robot companion. Perhaps there is a compromise available where the older person has the option of disrupting monitoring for short, or at least clearly defined, periods. This might appeal to older users who wanted to have very private conversations or engage in activity that they thought was very private. Consideration does, however, have to be given to a potential undesirable side-effect of this measure to protect privacy. This is that an older person may be coerced into turning off the monitoring facility by a carer wishing to conceal poor care (i.e. where a robot is also being used to police the standard of care provided). Mechanisms already exist for steps to be taken to protect vulnerable adults where there is suspicion of coercion, and regrettably many older people already experience sub-standard care and suffer abuse at the hands of their so-called carers that goes unreported.¹⁰ In this respect, an older person might not be worse off with a facility to disable monitoring than they would have been without the robot. A decision has to be made, therefore, about whether the potential policing capacity of the robot is sufficiently great to outweigh the threat to privacy of not being able to disable the monitoring facility.

Although many other matters could in principle be made subject to agreement, including the threshold that had to be reached for some monitored mishap to count as a genuine emergency, it is not necessary to go much further. The guiding thought is that the process for introducing the robot, as well as the robot itself, has to be sensitive to the wishes of the older person within certain limits. If they are not, both the design of the robot and its method of introduction into a household are ethically flawed.

Operationalizing autonomy is not only a matter of the agreement that lays down the ground rules. It is a matter also of what needs to be done to the robot in design terms. Clearly the robot has to be able to take 'No' for an answer – at least some of the time, and, depending on what the user has said 'No' to. Up to a point this lesson is already reflected in the current ACCOMPANY tablet interface, which both anticipates and accepts user non-co-operation, albeit at the "cost" to the user of being shown a sad face

¹⁰ O'Keeffe et al. reported that '2.6% of people aged 66 and over living in private households reported that they had experienced mistreatment involving a family member, close friend or care worker during the past year'. See O'Keeffe et al. 2007, Cooper et al. 2008.

by the tablet. The ACCOMPANY tablet design is also in keeping with the distinction between coercive persistence by the robot in enforcing an agreed care-regime, and the gentler suggestions with emotive masks of the current design. This may be as far as the current design can go in the direction of robot persuasion, which is clearly what the focus groups preferred to grinding persistence.

We can distinguish between two aspects of the ACCOMPANY tablet interface: its expressive, supposedly “empathic” capability, and its persuasive, but not overly insistent, prompting of the user. The second feature is clearly in line with the priority given in the value framework to autonomy. But empathic features are not, or are not necessarily, autonomy-producing. The reason is that the mask does not always show the robot entering into the emotions of the user, as genuine empathy would; instead, the sadness reflects a carer’s reaction to non-adherence with a care regime that is in the user’s interest. Similarly the tablet’s ‘happy mask’ sometimes expresses happiness at adherence, which is more paternalistic than autonomy-promoting. Sometimes the mask comes close to empathy in that it anticipates the user’s emotional response, albeit in a potentially infantilising way; for instance the positive, ‘happy mask’ prompted by the receipt at the door of a parcel for the user; or a negative, ‘sad mask’ when the user’s plant is knocked over.

There is scope for improving the tablet interface by confining uses of the mask to simulations of the user’s emotions in the course of actions by the robot that the user would like it to perform. The effect of this would be to abolish expressions of sadness for non-adherence. Another change that might promote autonomy is for the user to be given “ask-me-later” options on the interface. This would allow a user who was otherwise willing to comply with a demand of the care regime to time that adherence to suit him or herself. This option might not be shown where a medical need determined the timing of compliance; but it could be the autonomy-promoting default option.

5 Conclusions

The data we presented in D6.4 has provided insights into how older people and those who care for older people feel that the robot ought to be programmed to behave. Elsewhere we have published our thoughts on how the findings of D6.4 might:

- be used to improve robotic design (Draper et al 2014b);
- be applied to rehabilitation, including of social skills (Draper et al 2014a, Draper and Sorell 2014);
- help us better to understand how the robot might impact on the care triad (Jenkins and Draper 2014); and,
- support the hypothesis that assistive technology may undermine independence (Bedaf et al in progress).

In this deliverable we have reflected on how the data enriches our understanding of the values we identified in D6.2. The data were, however, prompted by specific scenarios. The scenarios were generated to emphasise potential tensions between the values of the framework, and whilst this achieved our aim of testing the value framework, the tensions that focused the scenarios will almost certainly have biased the results. More work will therefore be needed to get potential user groups to

engage with the ethical issues and underlying values, bearing in mind the difficulties of getting lay participants to engage with values in the abstract.

Our findings generally supported the weight given to autonomy, but suggested that safety issues may perhaps be more weighty than we had previously supposed. That said, the participants' concerns were subtle. The robot itself was not regarded as dangerous. Rather concerns seemed to centre on how safe it was to replace human judgement with robotic programming. Some of the concerns were highly paternalistic, which may reflect general attitudes to older people as well as concerns about the potential deficits of robotic care.

Our data echo the concern expressed more widely that robots should not be used to replace human-human interaction. Our data further suggest that human judgement cannot be replaced by robots. In this deliverable we have suggested how interactions with the robot could reflect the reciprocal care arrangements that older people value in other care/support arrangements. Mutual human-robot care will not, however, address the possibly increased social exclusion associated with robots as opposed to human-human care arrangements. Efforts must be made, therefore, to increase the range of interactions of users outside the home.

Our data suggest that a care-robot designed to be persuasive may be better than a robot designed to be persistent. Whilst the potential tirelessness of the robot overcomes the challenges to patience of human-human interaction, it can be associated with coercion, which acts against autonomy; persuasion, by contrast, facilitates autonomy. On the other hand, we have identified ways in which ostensibly persuasive techniques of robot-assisted care can be infantilising. Acceptable enablement is constrained by the need to change the behaviour of users in some cases whilst continuing to acknowledge their capacity for autonomous decision-making.

The perceived role of the robot is crucial to determining the norms against which the behaviour of the robot will be judged. The greater the variety of potential interactions between the older person and the robot, the greater is the potential for confusion about the appropriate norms to apply. This potential confusion may also encourage 'slippage' in that the older person – and others involved in supporting his or her independent living – may be inclined to manipulate the norms to de-emphasise enablement and independence. Devices simpler than companion robots might pre-empt this problem, but at the cost of eliminating a "presence" in the life of older persons (in the sense of "presence" outlined in D6.2).

Concerns about the potential of robots to erode privacy may extend beyond the user to the human-carers of that older person. Some formal carers raised the issue of the robot being used to 'spy' on them, whilst other formal carers did seem willing to use the robot to check up on, as well as to reinforce, adherence to treatment regimes. All forms of human care are likely to intrude to some extent on the privacy of the recipient of that care. Robots may be less intrusive by comparison. As for adherence, it does not seem acceptable to use the robot's data-recording capacities to second guess the older person's own testimonies. The value of the robot's capacity to retain and share information for the purposes of enablement is best maintained by ensuring that privacy norms are respected and the older person retains control of information that the robot gathers. Consideration of privacy in relation to multiple householders and issues of who 'owns' different types of information that the robot may collect were beyond the scope of our task in the description of work. Further work is undoubtedly

needed to classify the information that the robot collects and to establish criteria for legitimate access to and use of different kinds of information. This means taking account of the different kinds of value (commercial and ethical) of information the robot has to collect in order to maintain functioning. The aim of making the development of assistive technology profitable and affordable has to be set against the risks that older people will see no benefit from the value of the data generated about them. This leads us to our final conclusions in relation to the terms under which robots are introduced into the homes of older people.

We have signalled in several places the significance of achieving a shared understanding of the role, capabilities and potential behaviours of the robot. The values we have offered will need to be operationalized. One critical stage of operationalization is the introduction of a robot into someone's home. The value framework suggests that this should be a process rather than an event. We have demonstrated that agreements between providers and individuals have to be reached in order for tensions between the values in our framework to be resolved. These agreements cannot depend on generalised information about older users but need to be individualised. Having individualised agreements is in line with the invocation of prior agreements by participants when they tried to resolve tensions raised by the scenarios. Arriving at the right agreement depends on respecting the older person who is going to be subject to it, and ensuring that their autonomy and privacy are not considered less important than those of other adults. We have tended to suggest that autonomy overrides other values when there is a conflict, but it is not the only value relevant to care arrangement. Indeed, our participants thought other values, particularly safety, were sometimes nearly as weighty or equally weighty. Whether the value framework is operationalized effectively depends on older people being willing to exercise autonomy responsibly, and being helped to exercise autonomy responsibly. Operationalizing the value framework will help people to exercise their autonomy responsibly.

References

Age UK [online]. Available from: <http://www.ageuk.org.uk/get-involved/loneliness> [Accessed 19 Sep 2014]

Bedaf, S., Draper, H., Gelderblom, G-J., Sorell, T., de Witte, L. Can a service robot which supports independent living of older people disobey a command? The views of older people, informal carers and professional caregivers on the acceptability of robots. (In progress)

Berdichewsky, D., Neuenschwander, E.: Towards An Ethics of Persuasive Technology. *Communications of the ACM* 42(5), 51–58 (1999)

Birmingham Policy Commission. *Healthy Ageing in the 21st Century: the best is yet to come*. 2014 available: <http://www.birmingham.ac.uk/Documents/research/policycommission/healthy-ageing/Healthy-Ageing-Policy-Commission-Report.pdf>

Buetow, S., Kiata, L., Liew, T., Kenealy, T., Dovey, S. & Elwyn, G. Patient Error: A Preliminary Taxonomy. *Annals of Family Medicine*, 7 (3): 223-231 (2009)

Childress, J.F. (1982). *Who Should Decide?: Paternalism in Health Care*. Oxford: Oxford University Press

Cooper, C., Selwood, A. and Livingston G. The prevalence of elder abuse and neglect: a systematic review. *Age and Ageing*, 37 (2): 151-160 (2008)

Draper, H., Sorell, T. Patients' responsibilities in medical ethics. *Bioethics* (2002) 16(4): 335-353

Draper, H., Sorell, T. Telecare, remote monitoring and care. *Bioethics* (2013) 27(7): 365-372

Draper, H., Bedaf, S., Gutierrez-Ruiz, C, Lehmann, H., Hervé, M., Gelderblom, G-J., Dautenhahn, K., Amirabdollahian, F. What asking potential users about ethical values adds to our understanding of an ethical framework for social robots for older people. AIBS50 Congress 2014; Goldsmiths, London, UK (<http://doc.gold.ac.uk/aisb50/AISB50-S17/AISB50-S17-Draper-Paper.pdf>)

Draper, H., Sorell, T., Bedaf, S., Syrdal, D.S., Gutierrez-Ruiz, C, Duclos, A., Amirabdollahian, F. Ethical Dimensions of Human-Robot Interactions in the Care of Older People: Insights from 21 Focus Groups Convened in UK, France and the Netherlands. In: Beetz, M., Johnston, B., Williams, M.-A. (eds.) ICSR. LNCS (LNAI), vol. 8755, pp. 138–147. Springer, Heidelberg (2014)

Draper, H. and Sorell, T. Using robots to modify the demanding or impolite behavior of older people. In: Beetz, M., Johnston, B., Williams, M.-A. (eds.) ICSR. LNCS (LNAI), vol. 8755, pp. 126–135. Springer, Heidelberg (2014)

Elster, J. *Sour Grapes: Studies in the Subversion of Rationality*. Cambridge: Cambridge University Press (1983)

Ho, W.C. Dautenhahn, K., Burke, N., Saunders, J., Saez-Pons, J. (2013) Episodic memory visualization in robot companions providing a memory prosthesis for elderly users. Proceedings 12th European AAATE (Associated for the Advancement of Assistive Technology in Europe) Conference, Vilamoura, Algarve, Portugal, 19-22 September 2013

Jenkins, S. And Draper, H. Robots and the division of healthcare responsibilities in the homes of older people. In: Beetz, M., Johnston, B., Williams, M.-A. (eds.) ICSR. LNCS (LNAI), vol. 8755, pp. 177-186. Springer, Heidelberg (2014)

Nussbaum, M. *Women and Human Development: The Capabilities Approach*. Cambridge: Cambridge University Press (2000)

O’Keeffe, M., Hills, A., Doyle, M., McCreadie, C., Scholes, S., Constantine, R., Tinker, A., Manthorpe, J., Biggs, S. & Erens, B. UK Study of Abuse and Neglect of Older People: Prevalence Survey Report. London: Comic Relief & the Department of Health.

Powers, P. Persuasion and coercion: a critical review of philosophical and empirical approaches. *HEC Forum* (2007) 19 (2): 125–143.

Sharkey, A. Robots and human dignity: a consideration of the effects of robot care on the dignity of older people. *Ethics and Information Technology*, 16: 63-75 (2014)

Sorell, T., Draper, H. Robot carers, ethics and older people. *Ethics and Information Technology* (2014) 16: 183-195 DOI: 10.1007/s10676-014-9344-7 (March 27)

Spahn, A.: And Lead Us Not Into Persuasion..? Persuasive Technology and the Ethics of Communication. *Science and Engineering Ethics* 18, 633–650 (2012)

Walker, J.D. Liberalism, Consent, and the Problem of Adaptive Preferences. *Social Theory and Practice* 21 (3): 457-471 (1995)