A System for Debiasing the Excessive Weight of Momentary Encapsulation in Decision-Sensitive Situations

Amir Konigsberg
General Motors R&D
Email: amir.konigsberg@gm.com

Ron Asherov
Email: ron.asherov@gm.com

***Abstract*—We describe a system that provides what we call all things considered support to a user. The core feature of this system is that it finds a balance between the satisfaction of short term (local) preferences and the satisfaction of long term (global) preferences. By operating according to both local and global standards the system serves a debiasing function – it produces recommendations that bypass the common tendency that people have of granting excessive weight to utilities that relate to the short term. The novelty of this system is that for every decision it has to make it considers a user's interests all things considered; it incorporates that user's local interests as well as his global interests.**

***Index Terms*—preference management, debiasing, decision support systems, choice bias**

I. Introduction

In making decisions about actions and choices about items, people often take into account short term and long term considerations. People have preferences that relate to the short term and preferences that relate to the long term. By the short term we mean preferences relating to considerations that characterize the *local* environment of a particular situation – for instance a decision about which route to take on the way home from work; or whether to take dessert at a particular restaurant on a specific occasion. By the long term we mean preferences that relate to considerations that do not relate to any specific situation but rather reflect the user's preferences in general, from a *global* perspective, all things considered, distinct from the local considerations that may arise in any particular case. A user may for instance have a long term preference for quiet, scenic driving routes, or for skipping dessert except on special occasions. These kinds of preferences are often different to the preferences that people have on specific occasions, due to contextual considerations and other situational factors (There are various parallels to this distinction in the literature. Schelling [1], for instance, related to the tension between behaviors that people feel they should have and those they find themselves wanting; Others, namely Bazerman [2], suggest that there are two selves – a want self and a should self, and that each of these have competing preferences. Others, Shefrin and Thaler [3], relates to "doers and planners." See: [4].

For a variety of reasons short term and long term preferences often clash when people make decisions. The force of short term considerations in local situations frequently leads to the neglect or dismissal of long term preferences ([5]-[11]). A person that has decided to go on a diet on Tuesday may find himself tempted by a cake offered at a party on Thursday. And depending on how appealing the cake is, he may find himself acting locally against his global preference for staying away from fatty foods. By doing so he actively goes against his global preference in order to meet the local considerations ([12]-[14]). In cases such as these long term preferences may not presently be apparent to the user; they may momentarily be forgotten, or they may be pushed aside by more immediate and presently evident considerations. Moreover, it is not only long term preferences that can be dismissed, if long term preferences are held too rigidly, short term preferences may be discounted altogether, blocking the way for instances that warrant exceptions from long term preferences (imagine a case where you decide to go on a diet and a week or so later you're lucky enough to be invited to a dinner cooked by the best chefs in the world. In this case not having the flexibility allowing local considerations to influence your judgments seems irrational. See also: [15]). Both kinds of cases can lead users to sub-optimal decision making, and perhaps more importantly, to decision making that they later regret ([16], [17], [18], [8]).

The natural thing to do if you know that in some future situation you may not act as you presently hope you would is to commit to acting they way you want yourself to act. The hope is that if you *commit* to act in a certain way then you have the burden of commitment to deal with when the time comes, in addition to the utility considerations that normally guide you in choosing courses of action. Yet just like long term preferences and global considerations, commitments may also be overturned by considerations of the moment. Sometimes

Manuscript received November 17, 2013; revised February 10, 2014.

doi: 10.12720/joace.3.1.29-34

this will be justified, sometimes not. In the former case, specific evidence or the perceived value of present choices or actions may indeed merit overturning a previous commitment. In the latter case, acting again one's long term preferences and acting against one's commitment to honor these will lead to regret. Often this will lead to a twofold regret - for acting otherwise than one would have liked *as well as* regret for having breached a commitment, which can have a cost of its own, relative to upholding future commitments of the same sort [12][19]. What seems to be needed is a way of comparing the strength of long term preferences as well as the strength of short term preferences and short term utilities. Having such a mechanism in place can allow a user to balance out their actions rationally, both satisfying short term preferences and exercising the right amount self-control. It is a mechanism such as this that we propose here.

II. What We Want to do

To address this widely familiar problem we describe a system that provides what we call *all things considered* support to a user. The core feature of this system is that it finds a balance between the satisfaction of short term (local) preferences and the satisfaction of long term (global) preferences. By operating according to both local and global standards the system serves a debiasing function – it produces recommendations that bypass the common tendency that people have of granting excessive weight to utilities that relate to the short term. The novelty of this system is that for every decision it has to make it considers a user's interests *all things considered*; it incorporates that user's local interests as well as his global interests.

Broadly, the system learns the significance that a user attaches to each of his global preferences so as to represent how strongly he wants to maintain them, even when there is a pull from local considerations(Gul and Pesendorfer refer to "commitment rankings", which is somewhat similar to what we have in mind here. See: [20]. The system is then able to take into account pre-defined measures and considerations, and predict the long term effect (or cost) of a local decision. This prediction is better than a user's because the system always maintains a global perspective, and so it doesn't get encapsulated in the moment as a user frequently does. The main benefit is that the system produces recommendations that are better for the user in the long term while not ignoring the need to satisfy the demands of contingencies that may arise in the short term (while the system will take into account local considerations these may be neglected in the end result produced by the system. In such a case the system may simply find the global preferences stronger, in the present case, than local considerations). In doing so the system also induces trust in the user, and, more importantly, it procures support for actions that it believes the user will not regret later on.

III. Why DSS are Better Placed to Provide All Things Considered Decision Support

User-facing decision support systems typically perform one of two functions. They either *support* a user's decision making by, for instance, generating new evidence, revealing alternative and novel ways of addressing a problem, recommending actions or analyzing data, or they *substitute* the user by making his decisions for him [21] [22]. In both these capacities, decision support systems are preferred to human users in a number of senses.

First, the computational capabilities of a system enable quantitative analysis that can consider a much greater number of considerations, a great deal faster, and with greater precision, than a human user (additional advantages of decision support systems include: cost reduction, increased user satisfaction, improved effectiveness, time saving, facilitating communication between different people around specific problems, debiasing the decision making process and decreases errors). What is more, these systems are also expected to compute without any cognitive or computational biases that often distort the results produced by a human [23] [24] [25].

Second, computational systems, which represent a class of systems to which decision support systems belong, can consider both short-term and long-term goals *simultaneously*, in a non-biased way. What we mean by this is that there is no inherent computational limitation or technological barrier in terms of multiprocessing (the system's equivalent to a human's working memory, or attention), and therefore there is no *prima facie* reason to think that different kinds of decision-sensitive considerations – e.g., short term and long term goals, or preferences – cannot be taken into account with regard to the same action, at that same point in time.

Third, humans always make decisions from a particular point of view, in a particular, local decision environment. A user doesn't usually need to decide about whether or not to go to a movie *in general* but rather whether to go to a *specific* movie, on a particular day and time. And this local decision environment can impact decisions to the extent that considerations that the user has that are not presently apparent may be neglected. In utilitarian terms, the problem is that when a user is presented with two options, the option with the greater local utility (that is, utility with respect to local preferences) will in most cases have an advantage over the one with the higher global utility [4].

Fourth, for decision support systems there is no reason, cognitive or otherwise, that local considerations should impact decision making more than global considerations.

IV. The System

The system we propose debiases excessive weights granted to either local or global preference. Nonetheless in the present paper we focus cases in which the system can grant support to a user who grants excessive weights to local considerations.

As a first step, the system needs to know a user's global preferences. This has to be done before any specific situation where a decision is needed is encountered, so that the user gives the system his preferences from a neutral and clear point of view without the excessive utilities that often characterize local situations. This enables the system to provide a non-biased ranking, or weight, to the global preference. Later on, this will also enable the system to consider the global preference in a neutral, non-biased way, and in doing so permit it to bypass the problem of granting excessive utilities to local considerations.

The user is therefore asked to relay his preferences on various subjects in relation to different domains. For instance, within the domain of driving, the user may be asked about whether, in general, he prefers the fastest routes or the most scenic routes. Within the domain of tourism the user may be asked whether he prefers location to price when booking accommodation Next, the system asks the user to assign a weight to each global preference. This will be done by having the user provide a score to each preference. The score given to a preference will represent how strongly the user wants to satisfy this specific global preference, and will indicate the extent to which the user wishes to satisfy it even when granting excessive local utilities in specific situations. Once it has this information the system is able to support the user in local decisions when short term preferences are at play. The benefit of the system is that it considers the *a-priori* measures and considerations, and can thus predict the long term effect of a local decision in real-time better than the user who is often encapsulated by the situation itself.

In a local environment, when an actual local decision needs to be made, the system asks the user again about his preference, and these will now be tagged as local preferences. As before, the user provides the preference itself, as well as its weight (i.e. *right now* you want to get home as fast as possible). The system then considers every option in the option space, and examines it in light of *all* considerations, global and local. By combining all relevant utilities the system extracts a *total utility* of an option, which reflects its utility both in a global and a local sense. The best option would be the one with the maximal total (combined) utility. Fig. 1 (below) provides an outline of the system:

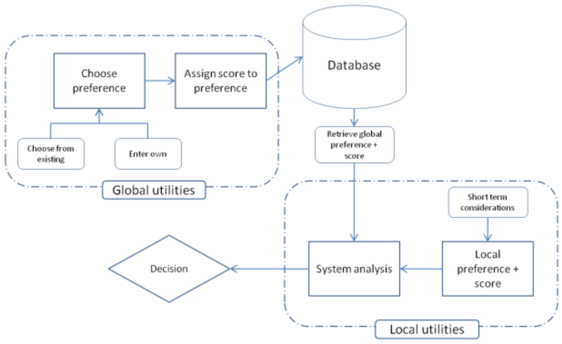

Figure 1. Outline of the system

V. COMPUTATIONAL OUTLINE

Above we provided a high level description of the system and its main operational principles. We will now outline the computational principles by which the system functions. As we have noted, in providing decision support, the system incorporates two types of considerations – global and local, corresponding to long term and short term preferences respectively. In order to be able to do this, the system needs to extract certain pieces of information:

- Firstly, information about specific preferences is needed. Namely, the system needs to know what the user wants to maximize in each case. For instance, the user may have a global preference for driving on scenic roads while often having local preferences to get home as fast as possible.
- Secondly, the system needs to know the extent, or weight (W_{global} and W_{local}), that represents *just how much* the user is eager to honor his preference in each case.
- Thirdly, the system needs to know the relevant utilities with respect to all related considerations, global and local. If there is a specific choice-option at hand, such as eating a cake, the system needs this option's utility with respect to the long term preference (which could be a diet) and the short term preference (which could be hunger or a craving for desert).

Once it has attained the values corresponding to each of these pieces of information, the system then proceeds to compute the total utility of an option or course of action, in light of both types of consideration, as a weighted sum:

$$U_{total} = W_{global} \cdot U_{global} + W_{local} \cdot U_{local}$$

A weighted sum is chosen because it allows the system to consider both global and local preferences in a way that considers the specifics of each global preference and local considerations. Even though you might normally prefer to refrain from eating sweets and fatty foods, when offered an exceptionally good cake, the expected enjoyment of the cake might exceed the cost of breaking your diet. And so as to be able to provide you with reliable and trustworthy decision support that is line with your best interests, *all things considered*, the system needs to know these specifics. Next, by comparing the utility values for each of the available options, the system advises you to take the option with the greatest total utility. This, in a nutshell, is the system. It is, we believe, simple yet effective, and most importantly, it caters to an as yet unmet need.

Let us now examine the behavior of the system by an example. Suppose you decide to go on a diet. When informing the system about your diet, it asks you how strong you intend for the diet to be. As it happens, you do want to lose weight, but you don't want your life to revolve around the diet fanatically, so you value your diet at 7 on a 1-10 scale. This score represents the global weight that you assign to your preference to go on diet. A few days later, a friend unexpectedly offers you a piece of a cake. Before deciding how to respond you decide to consult the system on whether to have the cake or not.

The system needs four values so as to be able to provide you with an all things considered recommendation:

A score that represents the local weight, which in turn represents the force of the local preference (the local weight represents how strong your local preference is, and this is independent of the local options at hand. Just like when you want to get home fast you have the same desire, be it via Road A or Road B).

- Continuing the previous cake example, this could mean how hungry you are or how strong your sweet tooth is at the present moment. For illustrative purposes let us suppose that you rank your hungriness at 4.
- A score that represents the global weight that you assign to maintaining your diet. This may also be understood as a score that represents how important it is for you to maintain your diet. As we assumed above, this score is 7.
- A score that represents local utility, which in our example can translate to how tasty the cake seems. Let us suppose that we perceive it as an 8.
- A score that represents the global utility, which corresponds to how the local option (eating the cake) satisfies the global preference (the diet). The user may not be aware of this factor. In this case global utility represents how much eating the cake will serve your diet. It obviously doesn't, so the effect is negative, for purposes of illustration let us say that it is a -5.

Plugging the values into the formula given above, the total utility of eating the cake is:

$$U_{total} = 7 \cdot (-5) + 4 \cdot 8 = \ -3 < 0$$

And as can be seen, in this instance we get a negative utility, and therefore a wise decision according to our system would be to reject our friend's offer of cake.

VI. Adaptation and Contextual Sensitivity

It is interesting to notice how our (informed) decision not to eat the cake can be reversed when the decision-relevant conditions change. Our system is in this sense adaptive, and this is another of its strengths; if scores change, then the system's recommendation might change too. The system is thus responsive to the user's context and to the decision environment. To illustrate this in relation to the cake example, there are several ways that the scores with which our system operates can change:

- If you do not immediately eat the cake but instead wait for an hour you may get hungrier, and so the score representing the local weight in the formula may increase, to 5 for instance. In this case the new total utility will be:

$$7 \cdot (-5) + \mathbf{5} \cdot 8 = \ 5 > 0$$

- A positive value here means that you eat the cake. Generally speaking, increasing W_{local} will make the local consideration more significant, hence if U_{local} is positive (as in our case, because the cake is tasty), the total utility will grow.
- Now suppose that your friend tells you that it's not just any cake that he's offering you, but your favorite rhubarb pie. We can envision that in such a case the score representing local utility jumps to 10, and the total utility to:

$$7 \cdot (-5) + 4 \cdot \mathbf{10} = \ 5 > 0$$

And here too, you eat the cake. As we can see, increasing the local utility U_{local} of an option increases the total utility U_{total} of an option (assuming the weight W_{local} is positive).

- Your friend tells you that it's a low-fat homemade cake, and so the negative impact that it has on your diet decreases to -4. In this case the total utility will be:

$$7 \cdot (\mathbf{-4}) + 4 \cdot 8 = \ 4 > 0$$

And here too you will accept your friend's offer of cake. Just like before, a change in an option's global utility affects its total utility.

It is important to notice that only three out of the factors can be altered in the local situation: these are the local value and weight, and the global value. As can be seen, no matter what happens, the constant global weight – W_{global} – of 7 remains constant. This is what prevents the force of local considerations and the weight of their corresponding utilities from biasing the choice.

VII. Implementation in Broader Option Spaces

Although this was an example of a yes/no question – to eat or not to eat the cake, the method we propose can be easily modified to fit any option space by comparing the total gain (utility) of the different options. In fact, one may claim that this is what was actually done above, with the second option being "not eating the cake," an option that received a total utility of zero (assuming both the local and the global utility of doing nothing is zero; There are a variety of differences between positing single yes/no options and positing those same options as separate options. For supporting research see: [26] [27]. Let us now demonstrate this modification of the system, from a question of whether to do something (eat the cake) or not, to a question of which option to take (what to eat, if at all), by expanding our option space to include two additional options – "cracker" and "fruit".

Assuming that crackers are healthier but fruits are sweeter, we conclude that the global utility of crackers will be higher than that of fruit and cake (because it serves your diet better). But we also assume that the local utility of fruit will surpass that of the cracker, because it is more tempting (more tasty perhaps). The results can be organized in Fig. 2:

	Nothing	Cake	Cracker	Fruit
Global weight	7	7	7	7
Global utility	0	-5	3	-1
Local weight	4	4	4	4
Local utility	0	8	-5	3
Total utility	0	-3	1	5

Figure 2. Table with results

The table shows the four factors that the system takes into consideration when computing the total utility of an option. The total utility of each of the options is given in the bottom row of the table. Now, being a maximizer, you choose the option with the greatest utility. Thus in the option space represented here you take the fruit.

Let us now briefly overview the range of possibilities within which each of the four factors needed to find the total utility remains constant. Firstly, so as to avoid self-serving biases, the global weight W_{global} cannot be changed in a specific situation in relation to local considerations. It therefore remains constant as long as it is relevant. In our case, the diet will be scored at 7 for the period of time you initially defined, and you cannot change your mind when food is at hand. If you want to change your global preference or its score, you need to do so at a point in time where there is no relevant specific option at hand. This is to avoid cases where local utilities can lead you to forfeit your previous commitment (This problem has been recognized in the literature and suggestions have been made about how to avoid breaking commitments because of local utility considerations while also enabling the flexibility of changing these commitments when one feels there is sufficient reason to do so. One such example are pre-determined monetary fines that a person can opt to. See: [28].

The local weight W_{local} can change from time to time, depending on the situation, but it cannot change during the time that options are being examined; as can be seen in the table above, the row corresponding to "local weight" remains constant, since the table represents a fixed point in time (a column that represents an option such as "an apple one hour from now" may have a different local weight). This prevents you from saying that you are hungrier when you see a cake (you are *justas* hungry as when you see the cracker; it's just that the cake is more appealing, and thus has a greater local utility). To conclude, the local weight remains constant for all options within a constant point of time – the point of time when you want to choose an option, namely the local decision environment – while the global weight remains constant for all options all the time, within the scope of the validity of the global consideration (i.e. until you actively and globally stop your diet under the restrictions discussed above).

VIII. Extensions of the System

One straightforward extension of the system is to include several preferences for both the global and the local kind, e.g. you are both on a diet, and trying to avoid sugar (two global preferences). In this case the system can simply sum the relevant weights and utilities to get this generalized form:

$$U_{total} = \sum W_{global} \cdot U_{global} + \sum W_{local} \cdot U_{local}$$

Another possible extension of the system will enable the system to learn the user's behavior in order to ensure less biased weights and more reliable recommendations as the system becomes more knowledgeable about the user. In this case the system can track the weights that the user assigns, together with his actual decisions in specific situations, and his *a-posteriori* thoughts (or regrets) about the choices he has made. Applying machine learning methods to this data can provide the system with insightful information about the user's behavior and the way he perceives conflicts between global and local considerations. And using this information we can further enrich the factors that contribute to the assessment of the total utility of an option:

$$U_{total} = C_g \cdot W_{global} \cdot U_{global} + C_l \cdot W_{local} \cdot U_{local}$$

With C_g, C_l being positive numbers that remain constant for all considerations at all times (until they are updated by analyzing new data) – that is, the constants remain the same when you are on a diet and when you want to get home fast. This extends the previous formula, since the previous formula can be written in the form above with $C_g = C_l = 1$. A simple way to understand the meaning of these constants can be demonstrated in the following way. Suppose that the system learns that despite the pre-defined weights and the system's recommendations, a user often decides to ignore global considerations and instead gives in fully to local utilities. Moreover, the system learns that even after some time has elapsed since the decision, the user does not regret his decisions, i.e. he still finds the (then) local temptations (cake) to be superior to the global consideration (diet), even after the specific situation has passed. A smart system that analyses these occasions can learn that the user tends to be too harsh with his global weights. And the system can therefore proactively decrease C_g to a lower value. The system will subsequently take local considerations into account to a greater extent than it had previously done. And perhaps on the next occasion it will advise the user to eat the cake. By following the system's recommendation to eat the cake, and later on not regretting having done so, the user can expect to develop an increasing degree of trust in the system.

A corresponding example in the opposite direction would be when, following the user's input regarding the global and local weights, tells the user to take the option with the greatest local utility (eat the cake); then, a few days (and pounds) later, the user regrets listening to the system's recommendation and eating the cake. The system may learn this pattern of behavior and increase C_g so that it will be less affected by local temptations, and advise the user to take options that he won't regret later on.

IX. Conclusion

We have outlined the main operational principles of a system that debiases user behavior in situations where global and local preferences can clash. The system supports decision making in real-time by providing a user with recommendations for action that balance global and local considerations by performing the necessary mutual tradeoffs based on solicited user-specific information. In doing so the system provides the user with all things considered decision support.

REFERENCES

[1] T. C. Schelling, *Choice and Consequence*, Harvard University Press, 1985.

[2] M. H. Bazerman, A. E. Tenbrunsel, and K. Wade-Benzoni, "Negotiating with yourself and losing: Making decisions with competing internal preferences," in *Academy of Management Review*, April 1998, vol. 23, no. 2, pp. 225.

[3] H. M. Shefrin and R. H. Thaler, "The behavioral life-cycle hypothesis," in *Economic Inquiry*, 1988, vol. 26, no. 4, pp. 609–643.

[4] K. L. Milkman, T. Rogers, and M. H. Bazerman, "I'll have the ice cream soon and the vegetables later: A study of online grocery purchases and order lead time," in *Marketing Letters*, Aug 2009, vol. 21, no. 1, pp. 17–35.

[5] R. Baumeister, *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*, NY: New York, 2003.

[6] G. Loewenstein and D. Prelec, "Anomalies in intertemporal choice: evidence and an interpretation," *The Quarterly Journal of Economics*, vol. 107, no. 2, pp. 573–597, May 1992.

[7] T. O'Donoghue and M. Rabin, "The economics of immediate gratification," *Journal of Behavioral Decision Making*, vol. 13, no. 2, pp. 233–250, 2000.

[8] G. Loewenstein and J. Elster, *Choice Over Time*, Russell Sage Foundation, 1992.

[9] C. F. Camerer, G. Loewenstein, and M. Rabin, *Advances in Behavioral Economics*, Princeton University Press, 2003.

[10] D. Kahneman, P. Slovic, and A. Tversky, *Judgment under Uncertainty: Heuristics and Biases*, 1st ed. Cambridge University Press, 1982.

[11] S. Plous, *The Psychology of Judgment and Decision Making*, 1st ed. McGraw-Hill, 1993.

[12] G. Ainslie, *Breakdown of Will*, 1st ed. Cambridge University Press, 2001.

[13] S. J. Hoch and G. F. Loewenstein, "Time-inconsistent preferences and consumer self-control," *Journal of Consumer Research*, vol. 17, no. 4, pp. 492–507, 1991.

[14] R. J. Herrnstein, G. F. Loewenstein, D. Prelec, and W. Vaughan, "Utility maximization and melioration: Internalities in individual choice," *Journal of Behavioral Decision Making*, vol. 6, no. 3, pp. 149–185, 1993.

[15] R. J. Aumann, "Rule-Rationality versus Act-Rationality," *Center for Rationality and Interactive Decision Theory*, Hebrew University, Jerusalem, Discussion Paper Series, 2008.

[16] M. Zeelenberg, "Anticipated regret, expected feedback and behavioral decision making," *Journal of Behavioral Decision Making*, vol. 12, no. 2, pp. 93–106, 1999.

[17] J. S. Dyer and R. K. Sarin, "Relative risk aversion," *Management Science*, vol. 28, no. 8, pp. 875–886, Aug 1982.

[18] J. Elster, *Ulysses and the Sirens: Studies in Rationality and Irrationality*, Rev Sub. Cambridge University Press, 1985.

[19] R. Nozick, *The Nature of Rationality*, New edition. Princeton University Press, 1994.

[20] F. Gul and W. Pesendorfer, "Temptation and self-control," in *Econometrica*, 2001, vol. 69, no. 6, pp. 1403–1435.

[21] F. Burstein and C. W. Holsapple, Eds., *Handbook on Decision Support Systems 1: Basic Themes*, 1st ed. Springer, 2008.

[22] J. P. Shim, M. Warkentin, J. F. Courtney, D. J. Power, R. Sharda, and C. Carlsson, "Past, present, and future of decision support technology," in *Decis. Support Syst.*, Jun. 2002, vol. 33, no. 2, pp. 111–126.

[23] J. Q. Chen and S. M. Lee, "An exploratory cognitive DSS for strategic decision making," in *Decision Support Systems*, Oct. 2003, vol. 36, no. 2, pp. 147–160.

[24] D. Arnott, "Cognitive biases and decision support systems development: a design science approach," *Information Systems Journal*, vol. 16, no. 1, pp. 55–78, 2006.

[25] H.-H. Huang, J. S.-C. Hsu, and C.-Y. Ku, "Understanding the role of computer-mediated counter-argument in countering confirmation bias," *Decision Support Systems*, vol. 53, no. 3, pp. 438–447, Jun. 2012.

[26] C. K. Hsee, "Less is better: When low-value options are valued more highly than high-value options," *Journal of Behavioral Decision Making*, vol. 11, no. 2, pp. 107–121, 1998.

[27] M. H. Bazerman, D. A. Moore, A. E. Tenbrunsel, K. A. Wade-Benzoni, and S. Blount, "Explaining how preferences change across joint versus separate evaluation," *Journal of Economic Behavior & Organization*, vol. 39, no. 1, pp. 41–58, May 1999.

[28] D. Ariely and K. Wertenbroch, "Procrastination, deadlines, and performance: Self-Control by precommitment," *Psychological Science*, vol. 13, no. 3, pp. 219–224, May 2002.

Dr. Amir Konigsberg, born in 1977, Netanya, Israel, is a senior scientist at General Motors Research and Development Labs, where he works on advanced technologies in the field of artificial intelligence and human machine interaction. Amir gained his PhD from the Center for the Study of Rationality and Interactive Decision Theory at the Hebrew University in Jerusalem and the Psychology Department at Princeton University.

Ron Asherov, born in 1994, is an intern researcher at General Motors Advanced Technical Center in Israel.