

DAVID GAL and BLAKELEY B. MCSHANE*

The question of how people should structure goal-directed activity to maximize the likelihood of goal attainment is one of theoretical and practical significance. In particular, should people begin by attempting relatively easy tasks or more difficult ones? How might these differing strategies affect the likelihood of completing the overarching goal? The authors examine this question in the context of an important goal for a large number of consumers—getting out of debt. Using a data set obtained from a debt settlement firm, they find that (1) closing debt accounts is predictive of debt elimination regardless of the dollar balance of the closed accounts, whereas (2) the dollar balance of closed accounts is not predictive of debt elimination when controlling for the fraction of accounts closed. These findings suggest that completing discrete subtasks might motivate consumers to persist in pursuit of a goal. The authors discuss implications for goal pursuit generally and for consumer debt management specifically.

Keywords: goal persistence, goal pursuit, subgoals, debt, consumer debt

Can Small Victories Help Win the War? Evidence from Consumer Debt Management

Successful goal pursuit often requires performing a set of actions (or subgoals) over an extended period of time (Anzai and Simon 1979; Newell and Simon 1972). For example, to write a scientific paper, a researcher might need to draft a conceptual development section, a methods and results section, and a discussion section. Similarly, to get out of debt, a consumer might need to pay off each of several individual debt accounts. In a broader sense, a person might group any number of tasks (e.g., errands) into a set or “checklist” of tasks that he aims to complete. Although the tasks composing this overall goal might be performed in parallel in some cases, in many other cases they will be—because of either necessity or convenience—performed in a temporal sequence of discrete steps (Newell and Simon 1972).

*David Gal (e-mail: d-gal@kellogg.northwestern.edu) and Blakeley B. McShane (e-mail: b-mcshane@kellogg.northwestern.edu) are Assistant Professors of Marketing, Kellogg School of Management, Northwestern University. The authors are grateful to Eric Anderson and to executives at Freedom Financial Network for their helpful comments. Both authors contributed equally to this research. James Bettman served as associate editor for this article.

As a result, the question of how people should structure the temporal sequence in which subgoals are pursued to maximize the likelihood of completing the overall goal is one of great theoretical and practical interest. For example, with respect to the preceding illustrations, should a researcher begin writing by addressing the complex conceptual development section of a paper first? Or should she begin with the relatively easier task of writing up the results? Likewise, should a consumer attempt to pay off small debts ahead of larger debts, or vice versa? More broadly, should people begin with the easy tasks on their to-do list so they can quickly check off a few items? Or should they begin with the more difficult tasks to get them out of the way?

Whereas past research in both marketing and psychology has identified important psychological processes that bear on this question (Amir and Ariely 2008; Anzai and Simon 1979; Bandura 1997; Catrambone 1998; Fishbach and Dhar 2005; Fishbach, Dhar, and Zhang 2006; Heath, Larrick, and Wu 1999; Khan and Dhar 2006; Shah and Kruglanski 2002; Thorndike 1898; Wertenbroch, Soman, and Chattopadhyay 2007), the mechanisms identified can lead to conflicting conclusions regarding whether, in practice, a person should begin goal pursuit by taking on relatively easy or relatively

difficult subtasks. For example, research indicates that the act of completing subtasks will increase self-efficacy (Bandura 1997), suggesting that people should complete relatively easy tasks first to experience the motivational benefits of greater self-efficacy. In contrast, other research shows that completing subgoals can lead people to switch their attention to unrelated goals (Fishbach and Dhar 2005; Kruglanski et al. 2002; Shah and Kruglanski 2002), suggesting that people might be better off maintaining their focus on an overall goal rather than on completing individual subgoals. Thus, although prior research on subgoal completion provides important theoretical insights, the sum of past findings does not readily yield actionable recommendations for goal pursuit in a real setting.

In the present investigation, we aim to provide insight into this question by examining how completing discrete subgoals influences overall goal attainment (beyond the absolute progress made toward the overall goal by virtue of subgoal completion) in the context of a long-term real-world goal. Specifically, we perform our investigation in the context of getting out of debt, a highly consequential goal for a large number of consumers. Two factors motivate our choice of this context. First, because most people are familiar with consumer debt and because consumer debts are generally spread out over many debt accounts, it provides a natural context in which to examine the effect of subgoal completion on overall goal attainment. Second, with global consumer debt measured in the trillions of dollars (Experian 2009), the best approach to reduce and eliminate debt has important implications for both consumers and policy makers.

To perform our investigation, we obtained a highly unique data set from a leading consumer debt settlement firm (viz., Freedom Financial Network). Using these data, we test whether closing individual debt accounts affects a consumer's likelihood of eliminating overall debt (regardless of the absolute amount of debt in the closed accounts). That is, at a particular point in a debt management program, is closing a higher fraction of outstanding accounts predictive of (eventually) closing all outstanding accounts when controlling for the fraction of total dollar debt contained in those accounts? The answer to this question should provide insight into the broader question of whether subgoal completion motivates people to persist in pursuit of a superordinate goal.

The remainder of our article is organized as follows. We next discuss prior literature in psychology and marketing pertinent to the question of how subgoal completion affects goal persistence. Subsequently, we describe the consumer debt settlement process and how its various features allow us to address our research question. Afterward, we describe our data set and the analyses we performed. We conclude by discussing theoretical and practical implications of our findings.

CONCEPTUAL DEVELOPMENT

The Psychology of Goal Progress

Prior research suggests that the perception of progress toward a goal has a strong and positive motivational effect on goal pursuit. In the context of choice, Soman and Shi (2003) show that, when choosing among alternative service routes (e.g., flights from New York to Paris), people prefer a path toward a goal in which progress is continuous rather

than a path that involves an interruption in forward progress. This holds even when both paths ultimately lead to reaching the goal in the same period of time (e.g., people prefer to take one step forward to taking one step back and two steps forward) and, in some cases, even when the interrupted path ultimately leads to more rapid goal attainment. Although Soman and Shi examine people's choice of paths toward a goal and not the temporal course of actual goal pursuit, their findings raise the possibility that consumers will be happier, and thus potentially more motivated to persist in pursuit of a goal, when they receive feedback that they are making progress toward their goal.

Consistent with this conjecture, research examining the temporal relationship between goal progress and goal persistence indicates that the perception of goal progress increases goal persistence. This finding, termed the "goal gradient hypothesis," has been observed among both animals (Brown 1948; Hull 1932, 1934; Miller 1944) and humans (Cheema and Bagchi 2011; Kivetz, Urminsky, and Zheng 2006; Nunes and Drèze 2006). For example, Kivetz, Urminsky, and Zheng (2006) find that coffee shop customers participating in a frequency reward program accelerated their purchases as they got closer to obtaining a reward.

Although a multitude of factors may be driving the relationship between the perception of progress and increased goal persistence, prior research has identified three main elements. First, a goal, by serving as a focal point of directed activity, is a natural reference point according to which people evaluate success or failure (Carver and Scheier 1998). As a result, the motivation to attain a goal is thought to mirror the pattern of prospect theory's value function (i.e., convex in the domain of losses and concave in the domain of gains [Kahneman and Tversky 1979]). Consequently, as a person moves closer to attaining his goal, the goal gradient is believed to become steeper, thus reflecting increased motivation to attain the goal. (Conversely, researchers believe motivation to engage in goal-congruent activity declines after the goal is attained, mirroring the concavity of the value function in the domain of gains [Heath, Larrick, and Wu 1999, Kivetz and Zheng 2006].) Second, the degree of perceived progress can serve to signal that the person is committed to the goal, thereby motivating accelerated goal pursuit (Fishbach and Dhar 2005; see also Staw 1981). Finally, goal progress can lead to increased goal persistence because the accomplished progress can bolster perceived self-efficacy with respect to the overall goal (Bandura 1997). These rationales and prior empirical work demonstrating that goal progress leads to increased goal persistence lead to our first hypothesis:

- H₁ (General): All else being equal, as the fraction of the distance to the goal that has been completed increases, so does a person's likelihood of attaining the overall goal.
- H₁ (Application): All else being equal, as the fraction of total dollar debt paid off increases, so does a person's likelihood of eliminating her debts.

The Psychology of Subgoal Completion

Objectively, subgoal completion is simply a marker of a certain amount of absolute progress. Therefore, one might expect that completing subgoals would have no impact on goal persistence independent of the effects of absolute progress (i.e., beyond the progress toward the over-

all goal implied by completion of the subgoal). However, research suggests several psychological factors associated with the mere completion of subgoals that might affect goal persistence.

Indeed, research shows that consumers frequently rely on discrete quantities as proxies for more continuous quantities even when the mapping between the discrete and absolute quantities is not especially meaningful (Amir and Ariely 2008; Gourville and Koehler 2004; Raghuram and Srivastava 2002; Shafir, Diamond, and Tversky 1997; Wertenbroch, Soman, and Chattopadhyay 2007). For example, Raghuram and Srivastava (2002) find that consumers tended to (1) spend less in a foreign currency than in their home currency when the foreign currency was nominally less valuable than the home currency and (2) spend more in the foreign currency when the foreign currency was nominally more valuable than the home currency. This suggests that consumers anchor to some degree on the number of currency units spent rather than solely on the purchasing power of the currency. Analogously, in the context of subtask completion, consumers' perceptions of progress are likely to depend on the number or fraction of subtasks completed rather than simply on the absolute progress made toward the goal. Thus, it could be expected that completing a discrete subtask would evoke a perception of progress and thereby motivate consumers to persist toward their overall goal (i.e., because perceived progress is thought to increase goal persistence through the routes discussed in the previous subsection, "The Psychology of Goal Progress").

However, the nature of subgoals suggests that completing them might give rise to additional effects on goal persistence beyond simply affecting the perception of progress. One reason is that subgoals, by allowing for natural breaks in the course of goal pursuit, often serve as the focus of goal-directed activity (Amir and Ariely 2008; Anzai and Simon 1979; Catrambone 1998; Newell and Simon 1972; Singley and Anderson 1989). As a consequence, a subgoal—rather than the superordinate goal—may become the most salient reference point for directed activity, thus resulting in increased motivation as a person gets closer to attaining the subgoal (i.e., the goal gradient becomes steeper as the person approaches the subgoal) followed by a subsequent decline. Concretely, this decline can be attributed to the idea that, by serving as the focus of directed activity, a subgoal can provide a tangible sense of achievement upon its completion.¹ While such attainment is intrinsically rewarding, the sense of achievement it provides can lead to a sense of complacency and thereby to reduced superordinate goal persistence (i.e., "resting on one's laurels"; Amir and Ariely 2008; Fishbach and Dhar 2005; Fishbach, Dhar, and Zhang 2006).

A focus on subgoal completion has another related consequence that is thought to arise because people frequently pursue multiple goals that compete for limited attentional resources (Kruglanski et al. 2002; Shah and Kruglanski

2002). Because the existence of a subgoal can shift a person's focus away from the superordinate goal and toward the subgoal, the sense of achievement that arises from successful completion of the subgoal might liberate or license the person to pursue an alternate—or even competing—goal rather than other subgoals associated with the original superordinate goal (Amir and Ariely 2008; Fishbach and Dhar 2005; Fishbach, Dhar, and Zhang 2006; see also Camerer et al. 1997; Heath, Larrick, and Wu 1999; Khan and Dhar 2006; Read, Loewenstein, and Rabin 1999). For example, Fishbach, Dhar, and Zhang (2006) find that participants who had completed an exercise subgoal on the path toward the superordinate goal of being fit were subsequently less likely to prefer eating healthy food than participants who did not complete the subgoal. Moreover, this effect was reversed when participants were explicitly cued with a reminder of their superordinate fitness goal. Thus, it appears that a focus on subgoal attainment can lead to neglect of the superordinate goal.

Although previous research has examined the negative motivational effects arising from the sense of achievement obtained by completing a subgoal, theory suggests that the sense of achievement attached to completing a discrete subgoal should also have positive motivational effects on goal pursuit. In particular, the law of effect (Thorndike 1898), which states that rewarded behaviors tend to be repeated, indirectly implies that the intrinsic reward associated with completing a subgoal might motivate people to continue goal pursuit. In this way, goal pursuit might become self-reinforcing, intensifying with the completion of additional subgoals. Thus, theory suggests that subgoal completion is likely to have both negative and positive motivational implications for superordinate goal persistence.

Temporal Dynamics of Subgoal Completion and Goal Persistence

Although research has identified both positive and negative motivational effects of subgoal completion on overall goal persistence, the practical significance of these findings for specific kinds of goals has not been extensively explored. In particular, many important goals (e.g., losing weight, getting out of debt, learning a trade) are pursued over long time horizons. Consequently, the motivational processes that are strongest in the immediate aftermath of subgoal completion might be less relevant for the pursuit of such goals than processes that persist or strengthen with time. Thus, understanding how the processes elicited by subgoal completion vary over time is important when structuring the pursuit of long-duration goals.

Some of the effects of subgoal completion on goal persistence are likely to be more transient than others (for summary analysis, see Table 1). In particular, the demotivational effects of subgoal completion arise from a focus on the subgoal and from a sense of achievement upon its completion; such effects are likely to be strongest in the immediate aftermath of subgoal completion and to then attenuate over time as attention presumably shifts back to the superordinate goal. Similarly, the positive reinforcement resulting from the intrinsic sense of reward associated with subgoal completion results from a focus on the completed subgoal and is thus similarly likely to be temporary.

¹Amir and Ariely (2008) use the term "achievement," whereas Fishbach, Dhar, and Zhang (2006) use the term "sense of progress" to refer to a construal of progress in terms of achievement. We use the term "sense of achievement" for this purpose to avoid confusion between a construal of progress in terms of achievement and other ways progress might be construed (e.g., in terms of commitment).

Table 1
TEMPORAL DYNAMICS OF MOTIVATION AFTER SUBGOAL COMPLETION

<i>Effect of Subgoal Completion on Superordinate Goal Persistence</i>	<i>Process</i>	<i>Posited Change with Increasing Temporal Distance from Subgoal Completion</i>
Complacency (decreased persistence)	Sense of achievement	Weakens
Switching to alternative goal (decreased persistence)	Sense of achievement	Weakens
	Subgoal directs attention away from superordinate goal	Weakens
Increased persistence	Positive reinforcement through sense of achievement	Weakens
	Self-efficacy arising from accomplished subgoals	Persists
	Completed subgoals elicit perception of progress leading to steeper goal gradient with respect to the superordinate goal	Strengthens
	Progress serves as signal of commitment to superordinate goal	Strengthens

In contrast, some processes by which subgoal completion increases goal persistence are likely to persist or even strengthen with increasing temporal distance from a completed subgoal. In particular, the positive motivational effect of increased self-efficacy evoked by completed subtasks is likely to be relatively stable because it reflects a person's belief in his ability to make progress toward—and ultimately attain—the superordinate goal. That is, a person's completed subgoals serve as testament to her ability to complete the superordinate goal both in the immediate aftermath of subgoal completion and subsequently.

Moreover, as a completed subgoal recedes into the past and attention shifts back to the superordinate goal, the superordinate goal should again become the more salient reference point. Consequently, over time, the more salient construal of the completed subgoal is likely to be that the person has made progress toward the superordinate goal rather than that the person has attained the subgoal (and can therefore rest on his or her laurels). This construal of the completed subgoal as representative of progress toward an overall goal can be expected to lead to increased superordinate goal persistence (i.e., as greater perceived progress leads to a steeper goal gradient).

Finally, as temporal distance increases from a completed subgoal, its completion should be more likely to be construed as a sign of commitment to the superordinate goal, thereby motivating increased goal persistence. This is because greater temporal distance from an event tends to focus attention on higher-order, relatively abstract aspects of the event (Liberman and Trope 1998; Trope and Liberman 2003). Given that goal commitment relates to the higher-level reason for goal pursuit (i.e., the “why”), people are

more likely to view distant versus proximate completed subgoals as indicative of commitment to their goal (Fishbach, Dhar, and Zhang 2006). Consistent with this reasoning, Fishbach, Dhar, and Zhang (2006) find in a vignette study that, when participants imagined making progress toward a goal in the distant (vs. proximate) future, they were more likely to focus on their commitment to the goal and more likely to express interest in pursuing additional activities related to the goal.

In summary, we surmise that the negative consequences of subgoal completion for superordinate goal persistence (and one of the positive consequences) are likely to be limited to the immediate aftermath of subgoal completion. Conversely, most of the positive motivational consequences of subgoal completion (i.e., increased self-efficacy with respect to the overall goal, a steeper goal gradient with respect to the overall goal, and increased commitment to the overall goal) are likely to persist or even strengthen over time. Thus, even though subgoal completion is often demotivating in the short run, we posit that the temporal dynamics of subgoal completion are likely to lead to a motivating effect on overall goal pursuit on a longer time scale. This discussion leads to our second hypothesis:

H₂ (General): All else being equal, as the fraction of discrete subgoals completed increases, so does a person's likelihood of attaining an overall goal.

H₂ (Application): All else being equal, as the fraction of total debt accounts paid off increases, so does a person's likelihood of eliminating his debts.

CONSUMER DEBT SETTLEMENT

Consumers who decide they want to reduce or eliminate their debts face a range of options. At one extreme, there are consumers who can afford to pay off their debts over time from savings and income without any changes to the structure of their debts; this strategy typically requires substantial lifestyle changes (i.e., expense reduction) on the part of the consumer. At the other extreme, there are consumers who choose to walk away from their debts, either by defaulting or by declaring bankruptcy. Between these extremes lie several alternatives.

Debt consolidation allows consumers to take out a single loan, which is then used pay off their other debt balances. The interest rate on this loan is typically lower than that on many if not all of the original balances, but, to obtain this lower interest rate, the loan must be secured by the consumer's assets, usually home equity. Thus, while secured debt consolidation can reduce a consumer's monthly payments, it can also put home equity at risk if she fails to meet the obligations.

Another alternative is credit counseling, which typically involves instituting a debt management plan. These plans are managed by a credit counselor and require consumers to make a single monthly payment to the counselor. The counselor subsequently disburses funds to creditors, typically at a lower monthly payment level than the consumer's previous payments (because the counselor is able to negotiate a reduced interest rate). While attractive, potential downsides of credit counseling involve (1) a long repayment period, (2) relatively small monthly payment reductions, and (3) a lower credit rating.

Debt settlement, also known as debt relief or credit advocacy, is another option, and it falls between credit counseling and bankruptcy in terms of the efficacy of eliminating debt and the negative effects on a consumer's credit. Designed for consumers who cannot afford to make monthly minimum payments to creditors, debt settlement plans have consumers stop making those payments and instead enlist the assistance of a debt settlement firm. The debt settlement program requires them to make a single monthly payment to a specially created personal savings account at a bank that serves as a dedicated account provider. The consumer designates that the bank give the debt settlement firm access to information regarding the transactions on the account (i.e., monthly deposits), and, equipped with this information, the debt settlement firm negotiates with the consumer's creditors to reduce the balance due on the consumer's debts. Reductions typically amount to approximately half the outstanding balance, which, in recent years, equates to approximately \$3,000–\$13,000 per customer across the entire customer base or \$8,000–\$19,000 per customer across those customers who successfully complete the program; the money saved in the account goes toward paying off these reduced balances. Overall, a debt settlement program typically takes three to four years until all debt is eliminated.

Several aspects of the debt settlement program make debt settlement data particularly suitable for studying whether closing accounts—regardless of their balance size—affects persistence in the goal of eliminating a consumer's debts. First, in addition to knowing the transactions associated with and the balance of the consumer's savings account, the debt settlement firm also has access to information regarding each of the consumer's outstanding credit card balances, thereby providing the firm with a picture of the consumer's overall progress toward eliminating debt. Second, the decisions regarding which accounts to negotiate and close at which times are made by the debt settlement firm rather than the consumer.² In deciding which accounts to negotiate and settle at a given time, the debt settlement firm takes account of several factors—namely (1) the relationships of the consumer's account representative (assigned by the debt settlement firm to settle a given client's accounts) with various creditors, (2) the ongoing negotiations of the debt settlement firm in general (i.e., independent of the individual account representative) with creditors over bulk settlement of multiple client accounts, and (3) the amount the consumer has saved to date. Because the decision to negotiate and close a particular account is made by the debt settlement firm rather than the consumer, potential biases are greatly reduced, thereby approximating a natural experiment (in which the treatment is the number or fraction of closed accounts and the outcome is the likelihood of total debt elimination). Third, consumers' monthly payment is a function of their total outstanding initial balance and, therefore, does not depend on the number of accounts they hold. This further eliminates biases and enhances the quasi-experimental nature of the setting. Fourth, virtually all accounts can be settled (i.e., small accounts are neither easier nor harder to settle than are large accounts), thus reducing

²The law requires that consumers approve and sign off on all settlements negotiated by the debt settlement firm. However, consumers refuse approval on less than .5% of proposed settlements so such refusals are negligible.

potential confounds with closing large versus small balances. Finally, in a debt settlement program (and in contrast to debt repayment schemes), accounts are either open or closed—there are no partial settlements—thus helping isolate the effect of closing accounts on consumers' goal persistence.

DATA

Our unique data set comes from a leading debt settlement company, which has settled debts on behalf of clients numbering in the hundreds of thousands. In particular, we have access to data from a random sample of 5943 clients who enrolled prior to January 1, 2007, and have since withdrawn, either successfully (i.e., paid off all balances in full) or unsuccessfully (i.e., left one or more balances not paid off in full).³ The principal data set consists of data from 4169 of the 5943 clients who successfully settle at least one account. Settling one account indicates that the debt settlement program is under way, and, as the data indicate, a non-trivial fraction of clients (i.e., $[5943 - 4169]/5943 = 30\%$) withdraw before the debt settlement firm has had the opportunity to negotiate even one debt on behalf of the client.

Table 2 presents a sample of our initial client-level data. For each of the 5943 clients, we know the date of enrollment, the total enrolled balance, and the number of accounts over which the enrolled balance was split. In addition, we know how the total debt was split over those accounts. For example, Client 2's \$12,602 initial debt was split over eight accounts each totaling \$500, \$3,739, \$800, \$1,731, \$2,596, \$1,795, \$736, and \$705.

In addition to this initial client-level data, we also have an event history for each of the 4169 clients for whom the program gets under way (e.g., for Client 2's data, see Table 3). This data set provides a running history of the settlements that the firm has negotiated on behalf of the client. For example, no settlement events occurred for Client 2 in the first three months. Client 2's first account, a balance of \$500, was settled during the fourth month after enrollment, and this balance was settled for \$250, one-half the initial amount. In the fifth month, two additional balances totaling \$2,531 were settled for \$1,567. Five more months passed without a settlement, and then two accounts were settled in the 11th month. By the 27th month, all eight of Client 2's balances had been settled, indicating successful completion of the

³Because it typically takes several years to successfully complete the debt settlement process, it is necessary to consider only clients who enrolled in years past. Including clients who enrolled in the present year or very recent years would bias the sample toward those who enrolled and dropped out unsuccessfully.

Table 2
INITIAL CLIENT INFORMATION

<i>Client Identification</i>	<i>Start Date</i>	<i>Balance (\$)</i>	<i>Accounts</i>
1	5/28/2006	15,804	6
2	6/28/2006	12,602	8
3	6/01/2006	23,498	2
4	5/28/2006	30,715	8
5	5/07/2006	17,330	5
⋮	⋮	⋮	⋮

Notes: This information is collected from the client on the day of enrollment and therefore exists for all 5943 clients.

Table 3
EVENT INFORMATION FOR CLIENT 2

Month	Settled Balance (\$)	Settled Accounts	Negotiated Amount
0	0	0	0
4	500	1	250
5	3,031	3	1,817
11	4,472	5	2,507
18	10,807	7	5,370
27	12,602	8	6,901

Notes: Settled balances, settled accounts, and negotiated amounts are cumulative from the day of enrollment. This information is collected throughout the debt settlement program and therefore exists only for the 4169 clients for whom the program gets under way.

program (i.e., if the number of settled accounts is less than the number of initial accounts in the terminal month of this client-level event history data, that indicates that the client withdrew from the debt settlement program unsuccessfully).

To provide an overall sense of the debt data, we present the distributions of the debt variables in Figure 1. As the figure indicates, all three groups of clients (i.e., those who never get under way, those who get under way but ultimately fail to eliminate their debt, and those who get under way and ultimately succeed in eliminating their debt) have similar distributions for the enrolled balance and enrolled accounts, with typical values being \$14,000–\$31,000 of total debt spread out over four to seven accounts.

When we consider the distributions of settled balances, settled accounts, and negotiated balances, the three groups—by construction—tend to differ. By definition, the “Never Under Way” group settles zero accounts and consequently has a zero settlement balance and a zero negotiated balance. Not surprisingly, the group that is ultimately successful tends to have larger values for these three variables than the group that ultimately fails. Finally, we note that the distributions of enrolled balance (accounts) and settled balance (accounts) for the group that is ultimately successful are, by definition, the same.

Using the data presented in Figure 1, we can thoroughly examine how settlements affect the likelihood of a success-

ful completion of the program. In particular, we can compare and contrast how the ratio of the dollar balance of the closed accounts to the total debt level—which represents the absolute progress toward the superordinate goal implied by completion of subgoals (i.e., the subject of H_1)—affects the probability of successfully completing the debt settlement program versus how the ratio of the number of closed accounts to the total number of debt accounts—which represents the fraction of discrete subgoals completed on the path toward the superordinate goal (i.e., the subject of H_2)—affects this probability.

METHODS AND RESULTS

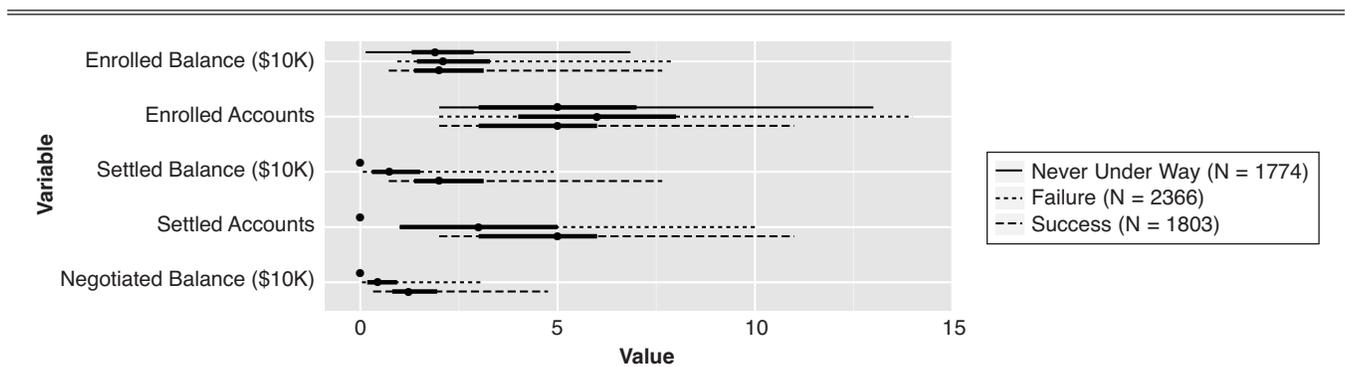
Debt Settlement and Identification

Imagine that two clients, Client A and Client B, enroll in the debt settlement program on the same day. Furthermore, imagine that these two clients are alike in every possible way. In particular, suppose that, upon enrollment, their enrolled balances are identical (e.g., they each enroll with one \$6,000 debt, one \$2,000 debt, and two \$1,000 debts). Using these two clients as exemplars, we discuss the debt settlement process and our strategy for estimating the role of the fraction of accounts paid on the likelihood of successfully completing the debt settlement program.

First, because Client A and Client B have the same total debt (i.e., \$10,000), each month they will deposit the same amount of money into their personal savings account at the bank that serves as a dedicated account provider. Consequently, both clients will be building up capital to pay down debts at the same rate. Now, suppose that after one year of enrollment, the debt settlement firm negotiates settlements for each client. In particular, suppose that the two \$1,000 balances of Client A are settled and the single \$2,000 balance of Client B is settled. Table 4 depicts such an outcome.

In this scenario, the two clients have identical debt profiles, and both have settled \$2,000/\$10,000 (or 20%) of their debt one year after enrollment. However, Client A has closed 2/4 (or 50%) of his accounts whereas Client B has closed only 1/4 (or 25%) of her accounts. The question of interest for us is whether Client A is more likely to successfully complete the program than Client B. In the context of

Figure 1
DISTRIBUTION OF VARIABLES BY CLIENT TYPE



Notes: The median of each variable is given by the dot, and 50% and 95% intervals are given by the thick and thin lines, respectively. The three groups have similar distributions of enrolled balances and enrolled accounts. For settled balances, settled accounts, and negotiated balances, the “Never Under Way” group is by definition fixed at zero; not surprisingly, the “Success” group has stochastically larger distributions for these three variables than the “Failure” group.

Table 4
BALANCE STATUS FOR TWO HYPOTHETICAL CLIENTS ONE YEAR AFTER ENROLLMENT

Client A		Client B	
Balance	Status	Balance	Status
\$6,000	Open	\$6,000	Open
\$2,000	Open	\$2,000	Settled
\$1,000	Settled	\$1,000	Open
\$1,000	Settled	\$1,000	Open

this pair, we can begin to answer this question by examining the data to determine whether (1) both Client A and Client B successfully completed the debt settlement program, (2) Client A successfully completed the program but Client B failed to do so, (3) Client A failed to complete the program but Client B did so successfully, or (4) both Client A and Client B failed to complete the program.

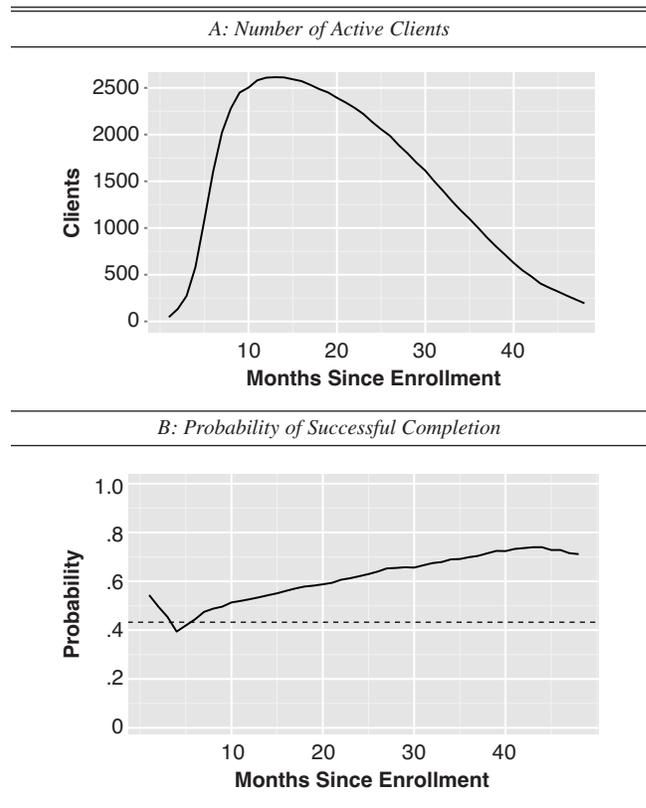
If we could construct perfect matched pairs for clients as we hypothetically outlined here, we could examine the frequency of each of the four outcomes. If the preponderance was like the second outcome, we could conclude that settling a greater fraction of the enrolled accounts (beyond the settled fraction of the enrolled dollar value of the accounts) leads to a greater likelihood of successfully completing the debt settlement program. In reality, it is impossible to construct such perfect matched pairs. However, using a regression model, we can approximate this matching process. Before turning to that, however, we first analyze our data in an exploratory manner.

Exploratory Data Analysis

We begin the discussion of our results by presenting some initial stylized facts about the data and then performing some preliminary analyses. We then delve into a more rigorous model-based analysis of the active clients. A client is defined as “active” in month $t > 0$ if (1) at least one account was settled for some $\tau \leq t$ and (2) at least one further account is settled for some $\tau > t$. This definition is important because it identifies the set of clients on behalf of which the debt settlement firm is actively and successfully negotiating. Clients who have not yet had one account settled (i.e., those who do not meet condition 1) are in the preliminary stages of the program; furthermore, for these clients, there is no variation between the dollar and number of settled balances, because both are zero. In contrast, clients with no settlements left (i.e., those who do not meet condition 2) are those who have either completed successfully or withdrawn unsuccessfully; as with those clients in the preliminary stages of the program, for clients in the former group, there is no variation between the dollar and number of settled balances because both are at the enrolled level.

In Figure 2, Panel A, we examine the number of active clients according to the time from enrollment. As the figure indicates, at any given time, no more than 2614 of the 4169 clients are active. Four years from enrollment, 95.3% of the clients in the database are inactive and, therefore, have either successfully completed the program or have withdrawn unsuccessfully (99.0% of the clients in our database have completed or withdrawn four years after their first settlement).

Figure 2
NUMBER OF ACTIVE CLIENTS AND PROBABILITY OF SUCCESSFUL COMPLETION



Notes: In Panel A, we plot the number of active clients by the number of months since enrollment. In Panel B, we plot the fraction of active clients who successfully complete the program by the number of months since enrollment; the horizontal dashed line gives the overall probability of successful completion.

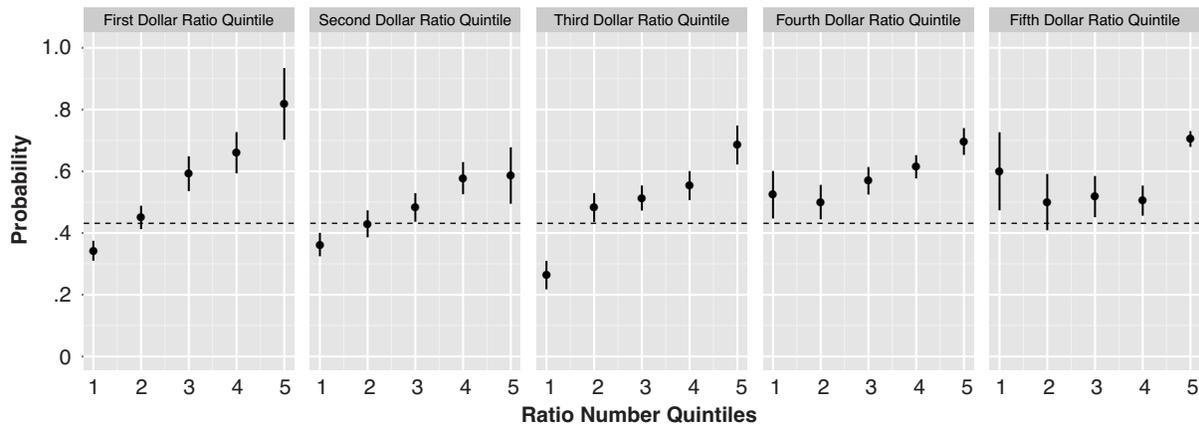
In Figure 2, Panel B, we determine the probability of successful completion according to the time from enrollment. The overall probability of successful completion is 43.2% (i.e., 43.2% of clients who get under way with the debt settlement program successfully complete it) and is indicated by the horizontal dashed line. As the figure shows, this probability steadily increases as a client remains active for more than four months. (The early decrease is likely due to small samples in the first few months and also that some clients withdraw in the first few weeks or months of the program because they realize it is not suited to them.)

We begin our analysis by investigating the data on the 2609 clients who are active one year after enrollment. Specifically, we are interested in whether the number ratio (defined as the number of accounts settled by time t [here, one year] divided by the number of accounts enrolled) can predict successful completion of the program. In particular, we are interested in how well this variable predicts successful completion compared to and even conditional on the dollar ratio (defined as the sum of the initial balances of the accounts settled by time t [here, one year] divided by total balance enrolled).

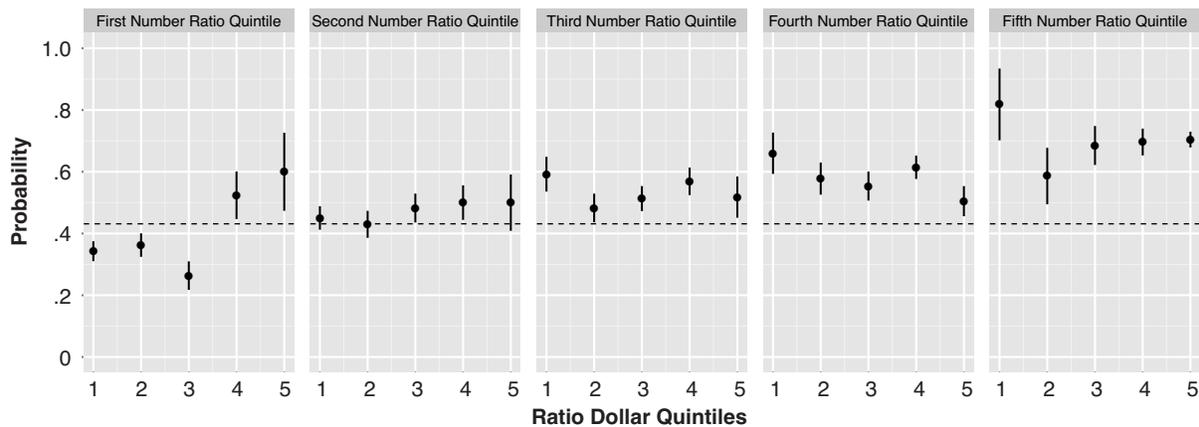
We conduct a preliminary analysis of this phenomenon in Figure 3, Panel A. For all clients who are active one year after enrollment, we calculate the number ratio and dollar

Figure 3
PROBABILITY OF SUCCESSFUL COMPLETION BY NUMBER AND DOLLAR RATIO QUINTILES

A: Probability of Successful Completion by Number Ratio Quintile Conditional on Dollar Ratio Quintile



B: Probability of Successful Completion by Dollar Ratio Quintile Conditional on Number Ratio Quintile



Notes: Consumers are placed in “bins” corresponding to the quintile of their number ratio and dollar ratio one year from enrollment. For each bin, the fraction of consumers who complete the program successfully is given by the dot, and plus or minus one standard error is represented by the vertical bar. The overall success rate of 43.2% is represented by the dashed horizontal line. For a fixed level of the dollar ratio quintile, an increase in the number ratio quintile is associated with increased probability of successful completion. For a fixed level of the number ratio quintile, the dollar ratio quintile has no clear association with probability of successful completion.

ratio at that time. Then, we group these clients according to quintiles of each variable and plot the probability of successful completion plus or minus one standard error for each of the 25 groups. The first graph in Figure 3, Panel A, shows that when we condition on the fraction of initial balance settled being in the bottom 20%, there is a substantial increase in the probability of successful completion as the fraction of the initial number of debts increases, dramatically rising from approximately 35% in the lowest group to approximately 80% in the highest group. A similar pattern generally holds across all graphs in Figure 3, Panel A: Conditional on the dollar ratio (quintile), an increase in the number ratio (quintile) correlates with a higher probability of successful completion.

In Figure 3, Panel B, we conduct the complementary analysis: We examine how the probability of successful completion varies when the dollar ratio quintile increases while the number ratio quintile is held fixed. As the figure demonstrates, there does not appear to be any clear associa-

tion within or across the bottom graphs. That is, conditional on the fraction of accounts paid off, the probability of success does not increase when a greater dollar percentage of debt is paid off.

Model

To formalize and extend the analysis presented here, we again work with data from the 2609 clients active one year from enrollment and examine how the dollar ratio and number ratio affect the probability of successfully completing the debt settlement program. In particular, we performed four logistic regressions: leaving both dollar ratio and number ratio out (the null model, Model 0), using each variable alone (Model 1D and Model 1N), and using both variables together (the “saturated” model, Model 2). Table 5 presents our results, which have several striking features. For example, both variables are highly statistically significant when analyzed on their own. When used jointly, however, only number ratio is statistically significant.

Table 5
LOGISTIC REGRESSIONS AT ONE YEAR

<i>Variable</i>	<i>Model 0</i>	<i>Model 1D</i>	<i>Model 1N</i>	<i>Model 2</i>
Intercept	.10*	-.28**	-.90**	-.91**
Dollar ratio		1.91**		-.25
Number ratio			3.21**	3.40**
Deviance	3609.4	3549.4	3474.2	3473.7

**p*-value between .01 and .05.

***p*-value between 0 and .001.

Notes: Values not marked with asterisks denote a *p*-value greater than .10.

We can formalize these results using the deviance of the various models (the deviance of a model is equal to $-2 \times \log$ -likelihood). When one model is a nested subcase of another model (e.g., Model 0 is a nested subcase of Model 1D because Model 0 is Model 1D with the parameter on dollar ratio fixed at zero), the deviance of the reduced model minus the deviance of the larger model is distributed as a χ^2_k random variable, where *k*, the degrees of freedom, is equal to the number of fixed parameters. Both Model 1D and Model 1N make large and highly statistically significant improvements over the null Model 0. That said, Model 1N constitutes a larger improvement: It reduces the deviance by 135.2 relative to Model 0, whereas Model 1D reduces the deviance by only 60.0 relative to Model 0.

The analysis presented in the previous paragraph examines each variable (i.e., dollar ratio and number ratio) entirely on its own. It is more important to consider how they jointly predict the probability of a successful completion and, in particular, the explanatory power of one variable conditional on the other variable being in the model. This is accomplished by comparing Model 1D and Model 1N with Model 2: When we add number ratio to the regression containing dollar ratio, the deviance improves dramatically, from 3549.4 to 3473.7. In contrast, when we add dollar ratio to the regression containing number ratio, there is a mere .5 improvement in the deviance (not statistically significant). This suggests that in the presence of dollar ratio, number ratio is a useful predictor; however, in the presence of number ratio, dollar ratio is not.

Thus far, we have examined the effect of dollar ratio and number on the probability of successful completion. However, our data set contains a much richer set of variables that might potentially affect the probability of success. We study the impact of these variables on our results in two stages, beginning with an initial set of covariates that are known on the day a client enrolls in the debt settlement program. These initial covariates serve to describe the initial debt load with which each client enrolls. In particular, because we have the dollar value of each debt for each client, we augment our models by considering five additional covariates: the natural logarithm of the total amount of initial debt, the number of accounts across which that debt is spread, the average debt (in thousands of dollars), the standard deviation of the debts (in thousands of dollars), and the entropy of the debts (like standard deviation, entropy is a measure of dispersion). By controlling for these covariates, which describe each client’s initial debt load, we more closely approximate the matching discussed in the “Debt Settlement and Identification” subsection.

We add these five variables to each of the four regressions from Table 5 and present the new results in Table 6. The initial covariates contribute a great deal to the explanatory power of the model, as is evident by comparing the deviance for a given model in Table 6 with its analogue in Table 5. Nonetheless, the pattern we observed in Table 5 holds for the models presented in Table 6: (1) Dollar ratio and number ratio each improve the null Model 0, (2) number ratio improves it substantially more than dollar ratio does, (3) dropping dollar ratio from the saturated Model 2 model has a small and statistically insignificant impact, and (4) dropping number ratio has a large and statistically significant impact. Again, in the presence of dollar ratio, number ratio is a useful predictor; however, in the presence of number ratio, dollar ratio is not.

As a further test, we added two more variables to the model, which we term our “contemporaneous debt covariates.” These variables are contemporaneous in the sense that they are not known on the date the client enrolls in the debt settlement program; rather, like dollar ratio and number ratio, they are only known *t* (here, 12) months from enrollment. The two variables we consider are the number of months since a client’s last settlement and the negotiated ratio (i.e., the total amount a client has paid for debts thus far settled divided by the sum of the initial balances of those debts). We included these variables because the recency of a successful settlement and the magnitude of successful settlements relative to the enrolled amounts could plausibly affect a person’s motivation to persist in a goal. Table 7 presents the results of these models one year from enrollment.

As with the initial debt variables, these contemporaneous variables add to the explanatory power of the model. Nonetheless, however, the now-typical pattern emerges: (1) Dollar ratio and number ratio each improve the null Model 0, (2) number ratio improves it more, (3) dropping dollar ratio from the saturated Model 2 model does not affect the model substantially or significantly, and (4) dropping number ratio does affect Model 2 substantially and significantly. Again, in the presence of dollar ratio, number ratio is a useful predictor; however, in the presence of number ratio, dollar ratio is not.

In summary, we conclude that our results support H_2 (i.e., all else being equal, as number ratio increases, so does the likelihood of completing the debt elimination program).

Table 6
LOGISTIC REGRESSIONS AT ONE YEAR WITH INITIAL COVARIATES

<i>Variable</i>	<i>Model 0</i>	<i>Model 1D</i>	<i>Model 1N</i>	<i>Model 2</i>
Intercept	1.09	.71	.35	.35
Log total debt	.05	.07	.01	.01
Number of debts	-.07	-.08	-.11*	-.11*
Average debt	.13**	.10*	.12**	.12**
Standard deviation of debt	-.12**	-.09*	-.12**	-.11**
Entropy of debt	-.88**	-.87*	-.39	-.41
Dollar ratio		1.84***		.12
Number ratio			2.64***	2.53***
Deviance	3382.1	3331.5	3300.7	3300.6

**p*-value between .01 and .05.

***p*-value between .001 and .01.

****p*-value between 0 and .001.

Table 7
LOGISTIC REGRESSIONS AT ONE YEAR WITH INITIAL AND
CONTEMPORANEOUS COVARIATES

Variable	Model 0	Model 1D	Model 1N	Model 2
Intercept	1.47	.50	.29	.31
Log total debt	.03	.06	-.04	-.04
Number of debts	-.06	-.07	-.10*	-.10*
Average debt	.12**	.09*	.12**	.12**
Standard deviation of debt	-.11**	-.08*	-.11**	-.11**
Entropy of debt	-.89**	-.87*	-.33	-.32
Time since last settlement	.02	.05*	.08***	.08***
Negotiated ratio	-.36	.15	-.08	-.10
Dollar ratio		2.00***		-.09
Number ratio			3.08***	3.15***
Deviance	3379.3	3325.6	3282.9	3282.9

**p*-value between .01 and .05.

***p*-value between .001 and .01.

****p*-value between 0 and .001.

However, they do not support H_1 (i.e., that, all else being equal, as dollar ratio increases, so does the likelihood of completing the debt elimination program). Although an increased dollar ratio does significantly correlate with an increased likelihood of successfully completing the debt settlement program when it is considered marginally as well as in the presence of our initial and contemporaneous covariates, when we controlled for number ratio, dollar ratio is not a significant predictor.

Robustness

While our findings are consistent under the three different covariate specifications shown in Tables 5–7, we subjected them to several additional tests. In this section, we briefly review several of them to further demonstrate the robustness of our results. First and foremost, the findings presented in the “Model” subsection consider a static picture, which includes only the 2609 clients who are active one year after enrollment. Consequently, idiosyncrasies pertaining to either (1) the time period one year after enrollment or (2) the particular 2609 active clients could be responsible for our results. Therefore, we more fully investigate the phenomena under consideration by using all possible time periods after enrollment for all 4169 clients.

In particular, we fit Model 0, Model 1D, Model 1N, and Model 2 (1) at every possible month after enrollment, (2) using the clients who are active at that month, and (3) using all three covariate specifications (i.e., no covariates, initial debt covariates, and initial and contemporaneous debt covariates). We compare the deviances of these models in Figure 4 (for coefficient estimates and 95% confidence intervals, see the Appendix). The solid and dashed curves represent the change in deviance as dollar ratio and number ratio, respectively, are added to the null Model 0 (top) or subtracted from the saturated Model 2 (bottom). As all six plots demonstrate, our findings are stable. First, each variable makes a highly statistically significant improvement to the null model, as indicated by both the solid and dashed curves in the top graphs appearing well above the 99% statistical significance threshold represented by the dash-dot line. Second, number ratio results in a markedly

larger improvement. Third, as the bottom graphs show, number ratio provides a dramatic and highly statistically significant improvement to models containing dollar ratio, but dollar ratio does not generally improve models that already contain number ratio. Fourth, these patterns hold regardless of the covariate specification. The consistency of these results is striking.

Second, we wanted to confirm that the results of Figure 4 were not sensitive to the assumption of linearity employed throughout. Therefore, we conducted two tests that allowed for nonlinearity. The first test involved replacing the linear dollar ratio and number ratio terms in our regression with binary variables indicating the quintiles of dollar ratio and number ratio. All results remained qualitatively the same. The second test for nonlinearities took a different approach. Rather than employing rough functions such as indicator functions of quantiles, we instead used smooth cubic B-splines (De Boor 1978; Hastie, Tibshirani, and Friedman 2009) with five degrees of freedom. Again, all results (i.e., those depicted in Figure 4) remained qualitatively the same.

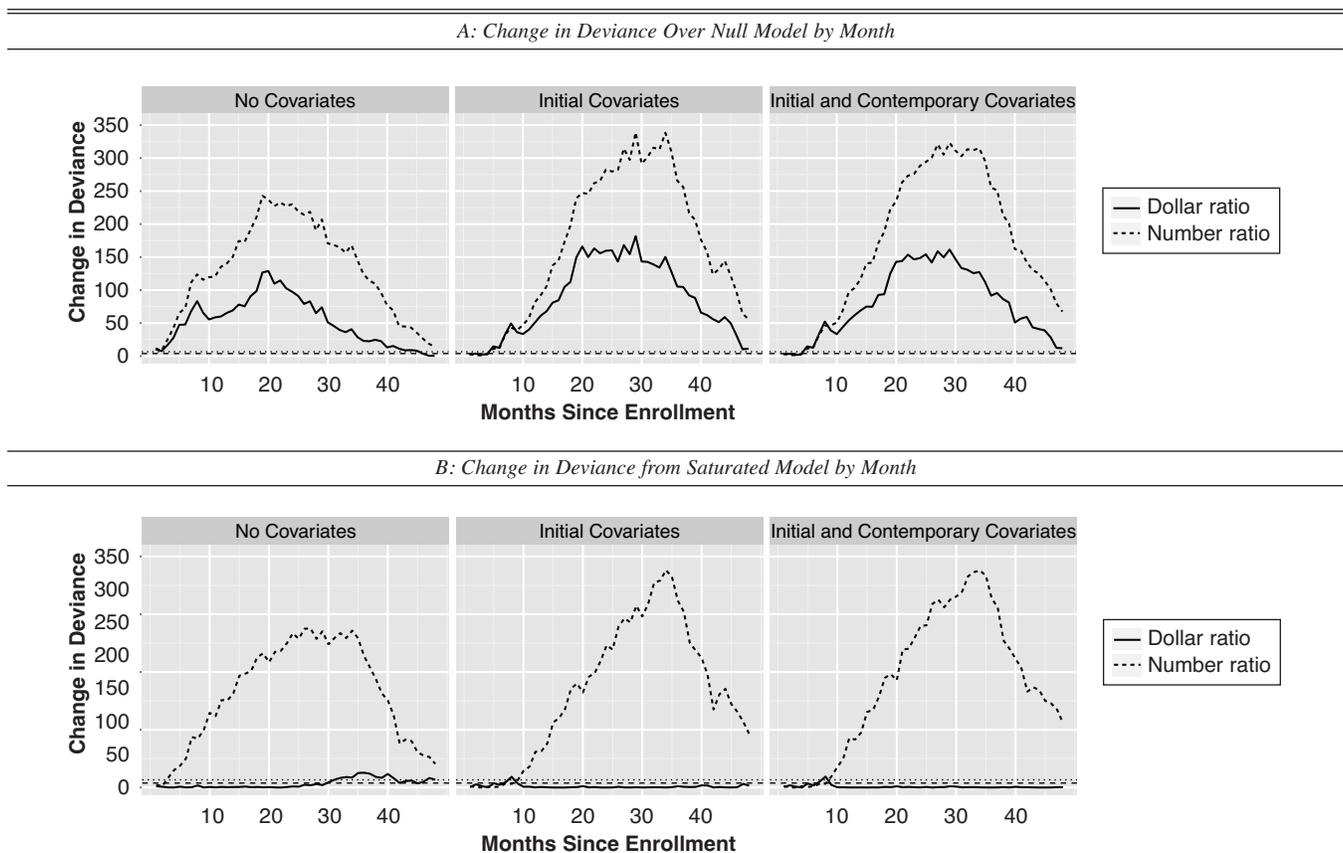
Third, we wanted to check whether our results were due to the definition of “active” we used. Consequently, we altered the definition such that a client is defined as “active” in month $t > 0$ if (1) at least one account was settled for some $\tau \leq t$ and (2) he has not successfully completed the program for some $\tau \leq t$. The difference in this definition and the original one is that it keeps people who have unsuccessfully withdrawn from the program as active (people who successfully withdraw become inactive upon debt completion, as before). Again, all qualitative conclusions remained the same.

The Starting Problem

As a final consideration, we examined whether our data had any bearing on a somewhat different question. Thus far, we have examined whether the completion of a subgoal (i.e., settling an account) leads to a greater likelihood of total goal completion (i.e., settling all accounts) regardless of the difficulty of the goal (i.e., balance size of the settled balance). Briefly, we consider whether the mere existence of a relatively achievable subgoal enhances the likelihood of early goal persistence.

People often have difficulty summoning motivation in the beginning stages of pursuit of a large goal, a notion termed “the starting problem” by Heath, Larrick, and Wu (1999; see also Locke et al. 1990). Researchers believe that one of the main benefits of dividing a goal into subgoals is that breaking a large, daunting task into smaller, relatively more manageable and more proximal tasks can promote goal initiation and persistence (Heath, Larrick, and Wu 1999; Sutton 2010; Weick 1984). For example, in a scenario study, Heath, Larrick, and Wu (1999) find that people believe a runner with a goal of running 1000 miles over three months is more likely to persist in the early stages of goal pursuit when the runner thinks of the goal in terms of running 11 miles a day than in terms of running 330 miles a month. Similarly, in an article titled “Small Wins,” Weick (1984) argues that large social problems are too enormous and overwhelming to tackle without breaking them down into more manageable steps. Likewise, Sutton (2010b) argues that “big, hairy, audacious goals” are often too large and

Figure 4
CHANGE IN DEVIANCE OVER NULL AND SATURATED MODELS BY MONTH



Notes: Panel A shows the change in deviance from the null model, and Panel B shows the change in deviance from the saturated model. The left graphs plot the change in deviance for models that use no covariates other than an intercept, dollar ratio, and number ratio; the middle graphs for models that also include initial debt level covariates; and the right graphs for models that include both initial and contemporaneous debt covariates. The change in deviance when dollar ratio (number ratio) is added to the model is represented by the solid (dashed) curve. When the curve is above the dashed (dash-dot) horizontal line, the change in deviance is statistically significant at the 95% (99%) level.

daunting to usefully guide practical action (see also Sutton 2010a). This discussion suggests the following hypothesis:

- H₃ (General): All else being equal, as the size of the smallest subgoal relative to the overall goal increases, a person's likelihood of initiating goal pursuit decreases.
- H₃ (Application): All else being equal, as the size of a person's smallest debt relative to her overall debt increases, her likelihood of settling one or more accounts decreases.

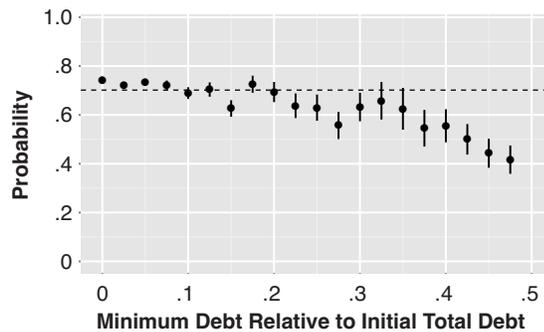
To examine this question, we consider the full data set of 5943 clients (i.e., the 4169 who settled one or more accounts and the 1774 who withdrew before settling a single account). We grouped clients by the size of their smallest debt relative to their total debt into bins of width 2.5% and computed the fraction who ever became active (i.e., ever settled a single account) and were therefore eligible to complete the program successfully. As Figure 5 indicates, the existence of a relatively achievable subgoal enhances the likelihood of settling an account: Those with large initial balances relative to total debt load are much less likely to become active and therefore succeed.

We test this hypothesis more precisely in Table 8. Having a larger minimum debt relative to total debt (i.e., a less achievable subgoal) leads to lower likelihood of becoming

active, as Model 1 shows. Model 2 presents a further test of this hypothesis, demonstrating that a larger minimum debt relative to total debt leads to lower likelihood of becoming active even when we control for our array of initial debt covariates. (Because we are modeling the probability of becoming active at time of enrollment, contemporaneous debt covariates, by definition, cannot exist.)

Although we have implicitly controlled for the dollar value of the minimum debt (it is included in total debt), we have not explicitly controlled for it, and it is possible that the dollar value of the minimum debt might have a different effect on the likelihood of becoming active than the total dollar value of debt. Moreover, it is possible that controlling for it would mitigate the statistically significant effect of the minimum debt relative to total debt. We test this possibility in Model 3, and indeed, the dollar value of the minimum debt has a statistically significantly different effect on the likelihood of becoming active from the effect of the total dollar value of debt. Nonetheless, the effect of minimum debt relative to total debt remains statistically significant and negative. Thus, we conclude that there is evidence that the mere existence of achievable subgoals motivates early goal persistence.

Figure 5
PROBABILITY OF BECOMING AN ACTIVE CLIENT BY
SMALLEST DEBT



Notes: Consumers are placed in “bins” corresponding to their smallest debt size divided by their total debt size, where each bin is of size 2.5%. For each bin, the fraction of consumers who ever become active (i.e., settle at least one debt) is shown by the dot and plus or minus one standard error is represented by the vertical bar. The overall rate of becoming active (70.1%) is represented by the dashed horizontal line. As consumers have smaller debts relative to their initial total debt, they have a greater likelihood of becoming active.

Table 8
LOGISTIC REGRESSIONS MODELING THE PROBABILITY OF
BECOMING ACTIVE

	Model 1	Model 2	Model 3
Intercept minimum	1.11***	-3.58***	-3.41***
Debt fraction	-2.38***	-1.23*	-2.32**
Log total debt		.48***	.32*
Number of debts		-.07*	-.03
Average debt		.00	.00†
Standard deviation of debt		.00	.00
Entropy of debt		.30	.17
Log minimum debt			.21*
Deviance	7156.1	7106.9	7102.1

†*p*-value between .05 and .10.

**p*-value between .01 and .05.

***p*-value between .001 and .01.

****p*-value between 0 and .001.

DISCUSSION

Summary

The principal goal of this research is to illuminate the effect of completing subgoals on superordinate goal persistence. We performed our investigation in a unique context involving real-world behavior over a time horizon measured in years and with high-stakes consequences for the people involved. Specifically, we examined the effect of closing debt accounts on a debt elimination goal. Our main finding is that closing off debt accounts is predictive of a person eliminating debts at any point during participation in a debt settlement program—even when controlling for dollar balances of the closed accounts. Indeed, the fraction of debt accounts paid off appears to be a better predictor of whether the consumer eliminates his debts than the fraction of the total dollar debt paid off, despite the latter being a relatively more objective measure of progress toward the debt elimination goal. More striking still, the dollar balance of closed debt accounts was not predictive of debt elimination when accounting for the fraction of debt accounts closed.

The Role of Goal Pursuit Context

Our findings highlight the complex interplay of psychological factors that influence goal persistence. Whereas some recent research has focused predominantly on the demotivational effects of subgoal completion, we found a robust positive effect of subgoal completion on goal persistence. Thus, it is worth speculating on what caused the divergence between our findings and those of prior research. We do so by considering differences in the contexts in which the investigations were performed.

A notable difference in context between our investigation and prior research on subgoal completion is the time horizon involved. We examined goal persistence over a long time horizon, whereas past research has focused on goal persistence in the immediate aftermath of subgoal completion. Thus, as we discussed in the section “Temporal Dynamics of Subgoal Completion and Goal Persistence,” the temporal dynamics of the processes through which subgoal completion affects superordinate goal persistence might explain the divergence between our findings and those of prior research.

However, there are other differences in context that could potentially explain the divergence in findings. One such difference is that prior research has typically been performed in contexts in which partial progress toward a goal still implies a degree of success. For example, a person might fail to reach the goal of completing a marathon or of losing 50 pounds but might nonetheless achieve some satisfaction from having run 20 miles or from having lost 30 pounds. In contrast, failing to complete a debt settlement program most often implies either bankruptcy or default; thus, partial progress might not be perceived as having value in and of itself. Therefore, in contrast to goals for which partial progress counts, all-or-none framing might not provide people with a license to disengage upon subgoal attainment (because nothing has truly been achieved until the ultimate goal is attained). Instead, in the context of all-or-none framing, subgoal completion might be more likely to be interpreted as a signal of commitment to the goal and might thereby increase motivation.

With respect to this account, Amir and Ariely (2008) examine the effects of discrete progress markers that functioned as “mere subgoals” (the attainment of which provided no value if participants did not also reach the ultimate goal). In this setting, the authors find diminished superordinate goal persistence (reflected by participants taking a longer time to reach the goal) when participants’ goals were divided into subgoals than when their goals were not divided into subgoals. Nonetheless, Amir and Ariely’s (2008) findings do not preclude the possibility that all-or-none framing contributed to our finding that completing discrete subgoals predicted goal attainment. A potentially important difference between our investigation and that of Amir and Ariely’s is that in their experiments, all participants reached the ultimate goal and differed only in how fast they reached it. It may be that people construe a completed subgoal as a signal of commitment to the superordinate goal under all-or-none framing when reaching the goal is relatively difficult and distant (as in our investigation) but not when ultimately reaching the goal is a foregone conclusion (in which case, a person’s commitment to the goal is less relevant).

In addition to potentially magnifying the effect of all-or-none framing on goal persistence, the relatively difficult and distant goal we investigated might have amplified the

importance of self-efficacy in motivating goal pursuit compared with investigations involving relatively simple and immediate goals. This is because people's belief that they are capable of achieving their goal is likely to be more relevant to goal attainment in the context of a difficult and distant goal than in the context of a simple and relatively immediate goal. Thus, to the extent that subgoal completion promotes superordinate goal persistence through enhancing self-efficacy, subgoal completion can be expected to be more effective in promoting superordinate goal persistence in the context of a relatively difficult and distant goal than in the context of a simple and relatively immediate goal.

Another potentially important distinction between our investigation and Amir and Ariely's (2008) prior work on subgoals is that they manipulated the number of subgoals into which their task was divided (finding that people tended to be slower to complete a task when it was divided into discrete subgoals), whereas we examined the effect of the fraction of subgoals that were completed. Indeed, in our main analysis, we controlled for the number of subgoals into which the overall goal was divided (i.e., we controlled for the initial number of debt accounts) while examining how the fraction of completed subgoals influenced goal persistence. Thus, it may be that completing subgoals per se does not lead to increased goal persistence but that increasing the fraction of completed subgoals (by serving as a proxy for progress) leads to increased goal persistence.

In this respect, a worthwhile question for further research might be to investigate how the number of subgoals into which a goal is divided affects goal persistence. Our finding with respect to the starting problem (i.e., that having an achievable subgoal predicts initial goal pursuit) suggests that breaking down a goal into smaller subgoals might increase goal persistence by creating a series of relatively achievable steps. Furthermore, it suggests that breaking a goal into a relatively large number of discrete subgoals will be particularly beneficial for the pursuit of more difficult goals (because the advantage of a series of relatively achievable steps is likely to be more relevant when goals are difficult). Conversely, it is imaginable that, beyond a point, dividing a task into a relatively large number of subgoals might lead to a demotivational effect on goal persistence because it may create the illusion of a relatively long distance to traverse to reach the overall goal. Moreover, if dividing a goal into a great many subgoals results in a relatively short distance between subgoals, a repeated demotivational effect in the immediate aftermath of each completed subgoal might overwhelm any long-term motivational boost of subgoal completion and may potentially lead to the diminished level of superordinate goal persistence that Amir and Ariely (2008) observe.⁴ Notably, in our model, we find a slight negative effect of number of debt accounts on attaining a debt elimination goal (see Tables 6 and 7); however, this result should be treated with particular caution because the number of debt accounts with which participants enter the program might be confounded with goal persistence through selection bias (e.g., those entering the program with a large number of debt accounts might be less responsible than those entering with a relatively small number of debt accounts).

In summary, contextual factors are likely to interact in complex ways to influence the course of goal pursuit. As a result, attempting to isolate processes through which subgoal completion affects goal persistence is likely to be insufficient to yield useful predictions regarding the effect of completing subgoals on goal persistence. Rather, analyzing how contextual factors such as time, all-or-none framing, task difficulty, and the number of subgoals moderate the relative strength of the psychological processes evoked by subgoal completion is necessary. Building on the work of others, we have performed some of this analysis in the present research and provided some tentative conjectures with respect to contextual moderators of the effect of subgoal completion on superordinate goal persistence that are consistent with our data. Nevertheless, much further research is needed to fully elucidate the effects of completing subgoals on motivation.

Practical Implications

Although we performed our examination in the context of a debt elimination goal, it is reasonable to speculate how our findings might apply to a broader set of contexts. In particular, our results lead directly to the hypothesis that in the context of pursuing a long-term goal, completing a greater share of goal-related tasks will tend to increase the likelihood of goal completion. If true, a possible recommendation arising from our findings is that when pursuing a long-term goal, people should focus on checking items off their list rather than focus simply on making progress toward their goal in an absolute sense (e.g., effort expended as a share of total effort required to complete the goal). In practice, this might imply focusing on completing relatively short and simple goal-related tasks ahead of relatively long and complex goal-related tasks. Notably, recent managerially oriented research offers a similar perspective at the organizational level, suggesting that managers can signal progress and competence through collective "quick wins" (Van Buren and Safferstone 2009). In a related vein, visual progress indicators, which are useful in motivating people to complete a goal (Cheema and Bagchi 2011), might be more effective in the context of a long-term goal when they display progress in terms of the share of tasks completed (e.g., a task checklist) rather than when they simply indicate the absolute progress made toward the goal.

Our findings also directly address how consumer psychology affects consumer behavior with respect to debt management. Consumers seem to believe that closing off debt accounts, regardless of balance size, is important in motivating them to persist in the goal of eliminating their debts (Amar et al. 2011). Likewise, the popular personal finance guru Dave Ramsey, while acknowledging that "the math" steers toward paying off higher-interest-rate accounts first, claims that his experience reveals that eliminating debts is "20 percent head knowledge and 80 percent behavior" and that people need "quick wins in order to stay pumped enough to get out of debt completely" (Ramsey 2009a; see also Ramsey 2009b). Our finding that closing off debt accounts—independent of the dollar balances of the closed accounts—is predictive of eliminating debts hints that this intuition has a basis in reality. In particular, our findings suggest that, consistent with the recommendations of financial advisors such as Ramsey, maintaining motivation to eliminate debts over a long time horizon might necessitate small wins along the way.

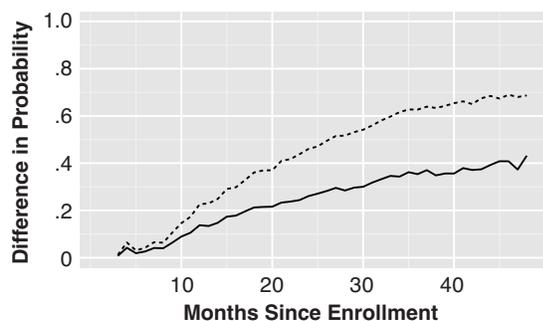
⁴It is also possible that, beyond a point, dividing a goal into a great many subgoals will lead to an evaporation of the distinction between discrete and continuous progress.

Indeed, this finding appears practically important for debt reduction, as we illustrate in Figure 6. Using our models estimated at each time period, we can simulate the chance that a client will complete the program under various debt settlement scenarios. In particular, holding the times of settlement and initial debt covariates constant, we consider three possibilities: (1) if the client were to always pay off the smallest remaining balance, (2) if the client were to always pay off a random remaining balance, and (3) if the client were to always pay off the largest remaining balance. The difference in the likelihood of successfully completing the debt settlement program under scenario 1 versus scenario 2 is illustrated by the solid line in Figure 6 and is substantial. For example, one year from enrollment, a client who has consistently paid down the smallest balances is 14% more likely to complete the program than one who paid down random balances. The difference is even more dramatic when we consider the smallest versus largest remaining balance strategies indicated by the dashed line.

Our findings also raise important policy questions. Currently, the U.S. government advises consumers to pay off their highest-interest-rate balances first (Federal Citizen Information Center 2011). Amar et al. (2011) make similar recommendations on the basis of normative principles. However, how can these normative prescriptions be reconciled with our findings suggesting a possible motivational boost acquired by closing debt accounts? Although it is undoubtedly true that when the interest rates and account balances of a given consumer's various debts diverge sufficiently, the benefits of engaging in the more rational strategy advocated by Amar et al. (2011) will outweigh the benefits of the extra motivation that is presumably acquired by closing debt accounts, the reverse might also often be the case. Perhaps the best option for policy makers is to simply inform consumers of both the rationally optimal approach to eliminate debt and the possible psychological benefits of closing account balances on persistence toward eliminating one's debts. Consumers can then make an informed decision.

Figure 6

SIMULATED DIFFERENCE IN PROBABILITY OF SUCCESSFUL COMPLETION



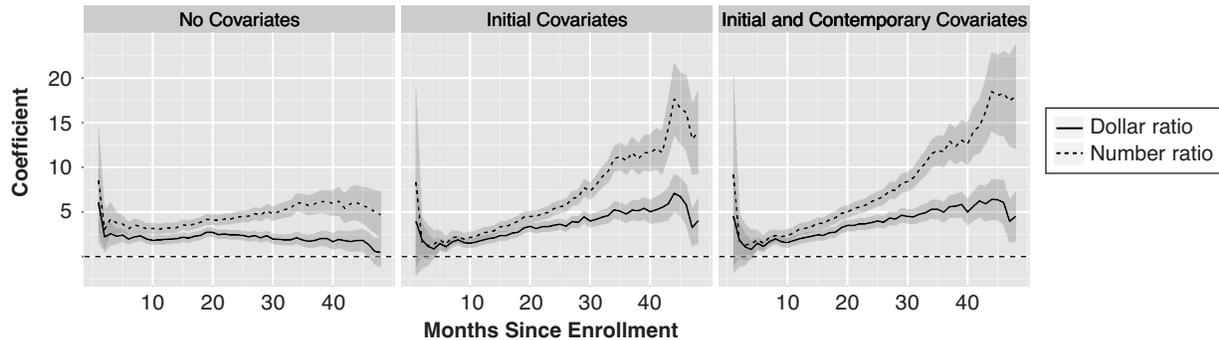
Notes: At each time period, the solid line shows the difference between probability of successful completion if clients always pay off the smallest remaining balance versus the probability of successful completion if clients pay off a random remaining balance. Similarly, the dashed line shows the difference between probability of successful completion if clients always pay off the smallest remaining balance versus the probability of successful completion if clients always pay off the largest remaining balance. Paying off the smallest remaining balance increases the probability of success.

REFERENCES

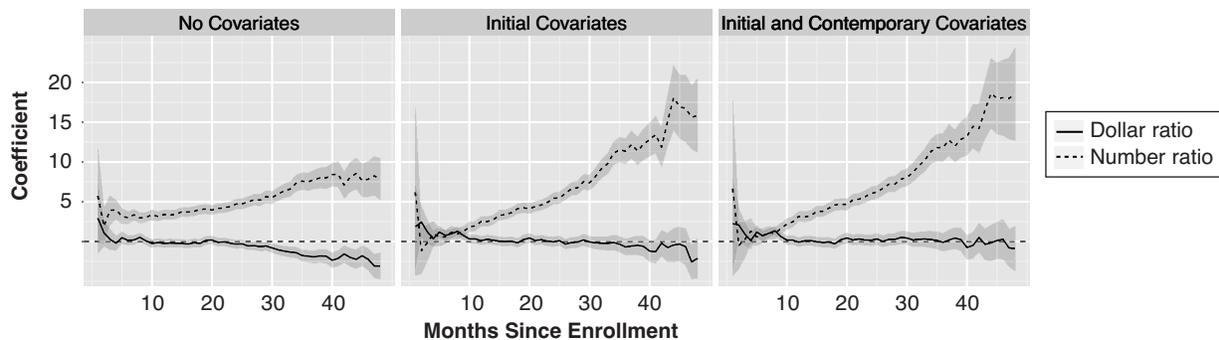
- Amar, Moty, Dan Ariely, Shahar Ayal, Cynthia Cryder, and Scott Rick (2011), "Winning the Battle but Losing the War: The Psychology of Debt Management," *Journal of Marketing Research*, 48 (Special Issue), S38–S50.
- Amir, On and Dan Ariely (2008), "Resting on Laurels: The Effects of Discrete Progress Markers as Subgoals on Task Performance and Preferences," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34 (5), 1158–71.
- Anzai, Yuichiro and Herbert A. Simon (1979), "The Theory of Learning by Doing," *Psychological Review*, 86 (2), 124–140.
- Bandura, Albert (1997), *Self-Efficacy: The Exercise of Control*. New York: W.H. Freeman.
- Brown, Judson S. (1948), "Gradients of Approach and Avoidance Responses and Their Relation to Level of Motivation," *Journal of Comparative and Physiological Psychology*, 41 (6), 450–65.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler (1997), "Labor Supply of New York City Taxi Drivers: One Day at a Time," *Quarterly Journal of Economics*, 112 (May), 407–441.
- Carver, Charles S. and Michael F. Scheier (1998), *On the Self-Regulation of Behavior*. New York: Cambridge University Press.
- Catrambone, Richard (1998), "The Subgoal Learning Model: Creating Better Examples so That Students Can Solve Novel Problems," *Journal of Experimental Psychology: General*, 127 (4), 355–76.
- Cheema, Amar and Rajesh Bagchi (2011), "The Effect of Goal Visualization on Goal Pursuit: Implications for Consumers and Managers," *Journal of Marketing*, 75 (March), 109–123.
- De Boor, Carl (1978), *A Practical Guide to Splines*. New York: Springer-Verlag.
- Experian (2009), "Marketing Insight Snapshot," (accessed April 24, 2012), [available at <http://www.experian.com/decision-analytics/market-insight-snapshot-june09-form.html>].
- Federal Citizen Information Center (2011), *Consumer Action Handbook*. Pueblo, CO: Federal Citizen Information Center.
- Fishbach, Ayelet and Ravi Dhar (2005), "Goals as Excuses or Guides: The Liberating Effect of Perceived Goal Progress on Choice," *Journal of Consumer Research*, 32 (3), 370–77.
- , ———, and Ying Zhang (2006), "Subgoals as Substitutes or Complements: The Role of Goal Accessibility," *Journal of Personality and Social Psychology*, 91 (2), 232–42.
- Gourville, John T. and Jonathan J. Koehler (2004), "Downsizing Prices Increases: A Greater Sensitivity to Price than to Quantity in Consumer Markets," HBS Marketing Research Paper Series, Harvard University.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2d ed. New York: Springer.
- Heath, Chip, Richard P. Larrick, and George Wu (1999), "Goals as Reference Points," *Cognitive Psychology*, 38 (1), 79–109.
- Hull, Clark L. (1932), "The Goal-Gradient Hypothesis and Maze Learning," *Psychological Review*, 39 (1), 25–43.
- (1934), "The Rats' Speed of Locomotion Gradient in the Approach to Food," *Journal of Comparative Psychology*, 17 (3), 393–422.
- Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47 (2), 263–92.
- Khan, Uzma and Ravi Dhar (2006), "Licensing Effect in Consumer Choice," *Journal of Marketing Research*, 43 (May), 259–66.
- Kivetz, Ran, Oleg Urminsky, and Yuhuang Zheng (2006), "The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusionary Goal Progress, and Customer Retention," *Journal of Marketing Research*, 43 (February), 39–58.
- Kruglanski, Arie W., James Y. Shah, Ayelet Fishbach, Ron Friedman, Woo Young Chun, and David Sleeth-Keppler (2002), "A Theory of Goal Systems," in *Advances in Experimental Social Psychology*, Vol. 34, M.P. Zanna, ed. San Diego: Academic Press, 331–76.
- Liberman, Nira and Yaacov Trope (1998), "The Role of Feasibility and Desirability Considerations in Near and Distant Future

Appendix
COEFFICIENT ESTIMATES COEFFICIENTS BY MONTH

A: Model 1D and Model 1N Coefficients by Month



B: Saturated Model Coefficients by Month



Notes: Panel A shows estimates from Model 1D and Model 1N, and Panel B shows estimates from the saturated model. The solid and dashed curves indicate the estimated coefficients for dollar ratio and number ratio, respectively; the shaded regions represent 95% confidence intervals. The left graphs plot the coefficients for models that use no covariates other than an intercept, dollar ratio, and number ratio; the middle graphs for models that also include initial debt level covariates; and the right graphs for models that include both initial and contemporaneous debt covariates.

- Decisions: A Test of Temporal Construal Theory,” *Journal of Personality and Social Psychology*, 75 (1), 5–18.
- Locke, Edwin A., Gary P. Latham, Ken J. Smith, and Robert E. Wood (1990), *A Theory of Goal Setting and Task Performance*. Englewood Cliffs, NJ: Prentice Hall.
- Miller, Neal E. (1944), “Experimental Studies of Conflict,” in *Personality and the Behavior Disorders*, Vol. 1. Oxford: Ronald Press, 431–465.
- Newell, Allen and Herbert A. Simon (1972), *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.
- Nunes, Joseph C. and Xavier Drèze (2006), “The Endowed Progress Effect: How Artificial Advancement Increases Effort,” *Journal of Consumer Research*, 32 (4), 504–512.
- Raghubir, Priya and Joydeep Srivastava (2002), “Effect of Face Value on Product Valuation in Foreign Currencies,” *Journal of Consumer Research*, 29 (3), 335–47.
- Ramsey, Dave (2009a), “Get out of Debt with the Debt Snowball Plan,” DaveRamsey.com, (August 1), (accessed April 24, 2012), [available at <http://www.daveramsey.com/article/get-out-of-debt-with-the-debt-snowball-plan>].
- (2009b), *The Total Money Makeover: A Proven Plan for Financial Fitness*. Nashville, TN: Thomas Nelson.
- Read, Daniel, George Loewenstein, Matthew Rabin (1999), “Choice Bracketing,” *Journal of Risk and Uncertainty*, 19 (1), 171–97.
- Shafir, Eldar, Peter Diamond, and Amos Tversky (1997), “Money Illusion,” *Quarterly Journal of Economics*, 112 (2), 341–74.
- Shah, James Y. and Arie W. Kruglanski (2002), “Priming Against Your Will: How Goal Pursuit Is Affected by Accessible Alternatives,” *Journal of Experimental Social Psychology*, 38 (4), 362–82.
- Singley, Mark K. and John R. Anderson (1989), *The Transfer of Cognitive Skill*. Cambridge, MA: Harvard University Press.
- Soman, Dilip and Mengzhi Shi (2003), “Virtual Progress: The Effect of Path Characteristics on Perceptions of Progress and Choice,” *Management Science*, 49 (9), 1229–50.
- Staw, Barry M. (1981), “The Escalation of Commitment to a Course of Action,” *Academy of Management Review*, 6 (4), 577–87.
- Sutton, Robert I. (2010a), *Good Boss, Bad Boss: How to Be the Best ... and Learn from the Worst*. New York: Business Plus.
- (2010b), “Hey Boss—Enough with the Big, Hairy Goals,” *HBR Blog*, (June 15), (accessed June 18, 2012), [available at http://blogs.hbr.org/sutton/2010/06/hey_boss_enough_with_the_big_h.html].
- Thorndike, Edward Lee (1898), *Animal Intelligence: An Experimental Study of the Associative Processes in Animals*. New York: Macmillan.
- Trope, Yaacov and Nira Liberman (2003), “Temporal Construal,” *Psychological Review*, 110 (3), 403–421.
- Van Buren, Mark E. and Todd Safferson (2009), “The Quick Wins Paradox,” *Harvard Business Review*, 116 (January), 54–61.
- Weick, Karl E. (1984), “Small Wins: Redefining the Scale of Social Problems,” *American Psychologist*, 39 (1), 40–49.
- Wertebroch, Klaus, Dilip Soman, and Amitava Chattopadhyay (2007), “On the Perceived Value of Money: The Reference Dependence of Currency Numerosity Effects,” *Journal of Consumer Research*, 34 (June), 1–10.

Copyright of Journal of Marketing Research (JMR) is the property of American Marketing Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.