

Konon, Alexander

Conference Paper

Career choice under uncertainty

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel -
Session: Occupational Choice and Job Mobility: Empirical Evidence, No. D18-V2

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Konon, Alexander (2016) : Career choice under uncertainty, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel - Session: Occupational Choice and Job Mobility: Empirical Evidence, No. D18-V2, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg

This Version is available at:

<https://hdl.handle.net/10419/145583>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Career choice under uncertainty

February 2016

Contemporary theoretical literature on occupational choice consists mostly of models that treat choice outcomes as either deterministic or risky. This paper proposes taking a slightly more realistic perspective by constructing a general occupational choice model on the basis of the assumption that outcomes are partially uncertain such that some reward distributions are unknown. The change in perspective yields some major advantages: Learning and career trajectories, which in general cannot be generated by models with deterministic or risky rewards, become a natural feature of the dynamic solution of sequential occupational choice problems. Furthermore, earnings-puzzle-like observations can be explained by sufficiently high uncertainty aversion, as uncertainty aversion has a significant impact on learning. In addition, central model predictions are consistent with data on relative choice frequencies and enterprise death rates.

JEL CLASSIFICATION: D81, D83, L26, C63

KEYWORDS: Occupational choice, uncertainty, earnings puzzle, Bayesian multi-armed bandits

1 Introduction

Modeling that explains career choices is a central focus of labor economics. In the literature, there are various occupational choice models (see, e.g., Lucas 1978; Kihlstrom & Laffont 1979; Holmes & Schmitz 1990; Lazear 2005; Roessler & Koellinger 2012) explaining why some individuals become wage workers and why some individuals decide to start an own business, thus, becoming entrepreneurs. However, most models (for instance, all of the aforementioned models) assume that once a decision for entrepreneurship or wage work is made, individuals stay in their selected occupation until the end of their careers. A consequence of the assumption that decisions are final is that the theoretical literature neglects mixed career trajectories, where careers can include spells in multiple occupations. Mixed careers—a mixture of spells in entrepreneurship, wage work, and unemployment—are, however, common in empirical data (see Ferber & Waldfogel 1998; Williams 2000; Arum & Müller 2004). As career patterns shift from long-term employment in one firm to greater mobility between firms and occupations (see Arthur 1994; Hall & Moss 1998; Valcour & Tolbert 2003; Granrose & Baccili 2006; Sullivan & Baruch 2009; Biemann, Zacher & Feldman 2012), mixed careers can be assumed to become even more common and important in the future. An additional empirical regularity that is hard to explain with most contemporary occupational choice models is the earnings puzzle: That is, the observation that a large share of entrepreneurs could earn more in dependent employment but they do not choose to do so (Hamilton 2000; Astebro & Chen 2014).

In this contribution, I argue that the inconsistency between theory and empirical evidence is significantly reduced if we alter one basic modeling assumption. All of the models mentioned so far are constructed under the assumption that occupational choices are made in a deterministic or risky choice environment.¹ However, it is more realistic to assume that the choice environment has at least some elements of uncertainty. A good example is entrepreneurship, as it is reasonable to assume that entrepreneurial incomes are not deterministic but rather draws from some distribution; an assumption that is consistent with risk and uncertainty but not with deterministic incomes. A risky entrepreneurial environment would imply that income distributions in entrepreneurship are perfectly known. This is unrealistic as a general assumption since the distribution of rewards from entrepreneurship depends on many factors individuals cannot fully anticipate. Innovative entrepreneurship is essentially all about trying something never done before, with the implication that even the distribution of rewards from entrepreneurship might be unknown.

A similar argument applies to the distribution of rewards from wage work. Wage workers do not have full control over their careers. The probability to be promoted or to lose one's job can depend on factors that are not fully transparent to workers, but these factors can be important for income distributions. Furthermore, many skills wage workers need are acquired by formal training in schools or universities in controlled settings. Formal training and real-world situations can be quite different.

¹There are few models not using one of the two assumptions that are discussed in Section 2.

For instance, the stress level in real-world situations can be much higher than in controlled settings. Consequently, some important skills affecting rewards in wage work can only be revealed by actually performing some tasks on the job. Antonovics & Golan (2012) show that occupational-choice patterns and wage growth are consistent with the assumption that jobs only gradually reveal information about unknown workers' skills. Unknown distributions do not require that there is no information at all. Usually, there is some prior information about occupational options and this information is updated as new information is obtained.

The model I construct in this paper is the first adaption of the multi-armed Bayesian bandit problem to general occupational choice. A bandit model is a stylized representation of a sequential decision problem where arms of a bandit machine represent some options with unknown payoff distributions. Individuals (gamblers) obtain information about the reward distribution of a particular option by selecting it and observing the outcome. Given some assumptions on the time horizon and discounting of rewards, a gambler has to decide which arm (occupational option) to select in every period and, especially, when to switch between arms conditional on new information (see Gittins, Glazebrook & Weber 2011).

An occupational choice model based on the bandit approach has a number of interesting features. First, with the help of the model uncertainty can be discussed in a systematic Bayesian framework. Second, assuming that some choice outcomes are uncertain automatically induces learning. As learning processes are paralleled by transitions from entrepreneurship to wage work and *vice versa* to obtain information, the model generates mixed careers including spells in different occupations. Third, introducing uncertainty aversion into the model is a simple approach to discuss learning efficiency—it is, for example, intuitively plausible that uncertainty-averse individuals avoid experimenting with occupations. Lastly, the bandit is the simplest theoretical device to depict the tradeoff between exploitation and exploration inherent in occupational choice, as the problem of finding the best occupational option (the exploration element) and generating a high-enough income (the exploitation element) must be solved simultaneously.

This paper concentrates on a general formulation of an occupational choice model where the success probability in an occupation is treated as a model primitive. It is also possible to formulate a model where success probabilities depend on some lower-level determinants. In effect, the general formulation can be combined with a less general static occupational choice model to yield a more refined dynamic occupational choice model. Yet, even the general formulation delivers interesting results. The main finding is that high uncertainty aversion can generate two effects of particular interest. (1) Highly uncertainty-averse individuals do not transition enough; and when they transition, they transition rather inefficiently compared to optimal behavior, which is in line with empirical observations. (2) In a setting with high uncertainty aversion, the number of entrepreneurs can be permanently too high compared to optimal learning. This effect is produced by a combination of misleading prior knowledge, characterized by overconfidence with respect to entrepreneurship, and sufficiently high uncertainty aversion reducing

learning efficiency. The second effect implies that the earnings puzzle has a simple and empirically credible explanation: entrepreneurial overconfidence combined with high uncertainty aversion.

As the bandit model allows for the derivation of predictions with respect to choice frequencies and entrepreneurial survival (or death) rates, I test two central model predictions using, *inter alia*, data from the Global Entrepreneurship Monitor, the World Bank, and research of Hofstede, Hofstede & Minkov (2010) on cultural dimensions. The first prediction is that the propensity to avoid the most uncertain option is increased by the level of uncertainty aversion. The second prediction, where the underlying effect is also responsible for the earnings puzzle, is that higher uncertainty aversion is associated with higher entrepreneurial survival rates, as increasing uncertainty aversion reduces learning efficiency and the probability to transition from one occupation to another. I find evidence for both predictions in data.

The remainder of the paper is organized as follows. In Section 2, I review previous research related to my model. Section 3 depicts the model's setup. In Section 4, I discuss decision rules. Section 5 gives an in-depth examination of occupational choice behavior in a partially uncertain decision environment. In Section 6, I show that two central predictions of the bandit model are consistent with data. Section 7 concludes. The Appendix provides proofs not included in the main text, additional results on empirical validity, and data characteristics.

2 Related research

In this paper, I combine, and partially extend, three strands of research. More specifically, I build upon research on occupational choice, research on multi-armed bandits, and research on uncertainty aversion.

OCCUPATIONAL CHOICE: THEORY AND EMPIRICS. I discuss three aspects of research on occupational choice: occupational choice theory, the earnings puzzle and its possible explanations, and empirical evidence on career trajectories.

Occupational choice models provide different sophisticated explanations for why individuals become entrepreneurs or wage workers. Lucas (1978) puts emphasis on entrepreneurial talent. Kihlstrom & Laffont (1979) show that entrepreneurs need to be more risk tolerant than wage workers. Holmes & Schmitz (1990) point out that entrepreneurs are opportunity seekers. Lazear (2005) argues that entrepreneurs are jacks-of-all-trades, while workers are specialists.² Roessler & Koellinger (2012) explain the emergence of entrepreneurs in a job matching model, where individuals who are best managed by themselves and not by others become entrepreneurs. All of the aforementioned models have two features. First, they operate with deterministic (for instance, Roessler & Koellinger 2012) or risky (for instance, Kihlstrom & Laffont 1979) rewards. Second, the models are not constructed to explain career trajectories.

²See Wagner (2006) and Stuetzer, Goethner & Cantner (2012) for empirical evidence on Lazear's predictions.

I contribute to the theory of occupational choice by introducing a new model operating with uncertain rewards. To be more precise, I assume that reward distributions of some occupational options are unknown but I also assume that individuals are Bayesian learners reducing uncertainty by selecting unknown options and updating information on the basis of observed rewards.

There are few models with learning components in the literature. For instance, in the model of MacDonald (1988) predictions on the distribution of rewards are derived from an information accumulation process. However, existing learning models do not answer an important question: What determines learning *efficiency*? In the model presented in this paper, learning efficiency is a function of the psychological variable uncertainty aversion. In this way, my model not only explains how individuals learn but also why different individuals can arrive at different conclusions given the same learning time and information, which is important to explain an observation like the earnings puzzle.

The earnings puzzle is based on the counterintuitive observation that a large share of entrepreneurs could earn more in wage work (Hamilton 2000; Astebro & Chen 2014), raising the question of why entrepreneurship is selected when wage work generates higher payoffs. Moreover, entrepreneur cohorts tend to *permanently* earn less than employees (Astebro & Chen 2014), such that income differences are not reduced over time. To my best knowledge, the only model with learning components able to replicate the earnings puzzle is Astebro, Chen & Thompson (2011). However, the model of Astebro et al. (2011) has a significant weakness as according to it the only reason to start an own business is the inability of the market to frictionlessly assign workers to firms or workers to tasks. If the assignment mechanism would be frictionless, no individual would have an incentive to transition to entrepreneurship. Consequently, entrepreneurs generated by the Astebro et al. (2011) model are rather of limited interest as the model only generates necessity entrepreneurs, selecting entrepreneurship because of the lack of other good options, and ignores other types such as opportunity entrepreneurs.³

An alternative way to explain the earnings puzzle are non-monetary rewards. Blanchflower (2000), Hundley (2001), Benz & Frey (2008a), and Benz & Frey (2008b) demonstrate that entrepreneurs experience greater job satisfaction than wage workers. Thus, the earnings puzzle might be explained by permanent non-pecuniary gains, such as overall job satisfaction, that are higher for entrepreneurs than for wage workers but are not reflected by purely monetary variables such as earnings.⁴ However, no occupational choice model known to the author accounts for non-monetary rewards. This is an important gap in the theory on occupational choice that needs to be addressed.

In this paper, I argue that a simple explanation for the earnings puzzle, very much overlooked in the literature, is that entrepreneurs are overconfident (Busenitz &

³In addition, the Astebro et al. (2011) model is not “fully closed” because it does not provide an explanation for the origin of firms if assignments are frictionless.

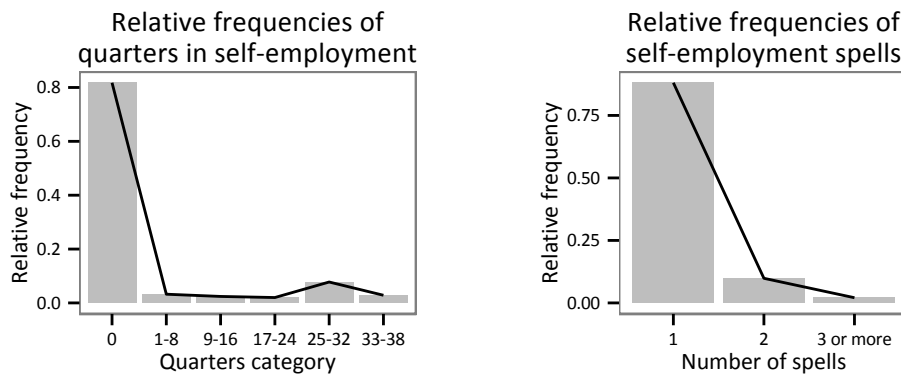
⁴There are also alternative explanations like a systematic under-reporting of entrepreneurial incomes by the self-employed (see Astebro & Chen 2014).

Barney 1997; Camerer & Lovallo 1999; Bernardo & Welch 2001; Koellinger, Minniti & Schade 2007) and uncertainty-averse (see Hofstede et al. 2010). The synergy between entrepreneurial overconfidence and uncertainty aversion plays a central role in the explanation of the earnings puzzle in my model.

Occupational choices are not final decisions. Burke, FitzRoy & Nolan (2008) show that pure entrepreneurial careers—an individual is an entrepreneurs over her whole career—are outnumbered by mixed careers—individuals combine entrepreneurship and wage work spells. In the same vein, Williams (2000) finds that a substantial share of US employees has previous self-employment experience (also, see Ferber & Waldfogel 1998; Arum & Müller 2004).

Figure 1a shows the relative frequency of quarters (three months) in self-employment on the basis of calculations provided by Burke et al. (2008), drawing data from the UK National Child Development Study for the years 1991–2000. Individ-

FIGURE 1: Observations on self-employment spells



(a) Relative frequencies of time spent in self-employment

(b) Relative frequencies of self-employment spells conditional on at least one spell

SOURCE: See Burke et al. (2008), especially Table 1 and 2 based on UK National Child Development Study (data for the years 1991–1999/2000).

uals with zero quarters in self-employment are “die-hard wage workers” who did not have an own business in the given time frame. The category of 33–38 quarters in entrepreneurship represents “die-hard entrepreneurs” who were mostly self-employed in the given time frame. As we can see in Figure 1a, there is a significant share of individuals of approx. 10% who are not in any “die-hard category,” i.e., they are neither pure wage workers nor pure entrepreneurs.

Figure 1b depicts the relative frequency of self-employment spells conditional on at least one spell (the data source is the same as in Figure 1a). The number of self-employment spells is a proxy for transitions between occupations: To generate a new self-employment spell, an individual must change occupations for at least a short period of time.⁵ According to Figure 1b, the overwhelming majority of entrepreneurs generates only one spell. Yet, approx. 12% generate more than one spell.

⁵Note, however, that some entrepreneurs, for instance, serial entrepreneurs, might start new busi-

Using Figure 1, I draw the following conclusion: Although mixed careers are relatively common and should be accounted for, the observable number of transitions between occupations is relatively low.

MULTI-ARMED BANDITS. The multi-armed bandit problem, as described in the introduction, provides a particularly suitable method to discuss uncertainty in occupational choice. Economic theory is aware of the bandit problem but there are only few adaptations. Rothschild (1974) uses a two-armed bandit to demonstrate that price distributions can exist even when costs and demand are identical across firms and markets. Jovanovic (1979) shows that employee turnover can be an unavoidable consequence of learning processes. Bergemann & Hege (2005) employ bandits to examine best financing rules for research projects with unknown success outcomes and unknown length. My contribution to research on multi-armed bandits is to apply the bandit approach to general occupational choice. This approach allows for the derivation of results on transition and learning efficiency not obtainable in occupational choice models without uncertain components.

UNCERTAINTY AVERSION AND ITS IMPACT. The attitude towards uncertainty is the central variable of my model. To avoid a model operating in an empirical vacuum, it might be convenient to scrutinize empirically observed uncertainty preferences. I concentrate on the cross-cultural dimension since it is better researched than intra-cultural differences in uncertainty preferences. The most prolific study evaluating cross-cultural differences in uncertainty avoiding behavior was conducted by Hofstede et al. (2010). Hofstede et al. (2010) assess uncertainty-avoiding behavior with the assistance of the following three questions answered by IBM employees:^{6,7}

- (Q1) “How often do you feel nervous or tense at work?”
- (Q2) The level of agreement with the following statement: “Company rules should not be broken—even when the employee thinks it is in the company’s best interest.”
- (Q3) “How long do you think you will continue working for IBM?”

According to Hofstede et al. (2010), a society has a high level of uncertainty aversion if members of this society often feel nervous at work (answer to [Q1]), society members want rules to be respected (answer to [Q2]), and members of such a society are reluctant to change their employer (answer to [Q3]). Hofstede et al. (2010) construct an index capturing attitudes towards uncertainty based on country-specific mean scores provided by answers to question (Q1)–(Q3).

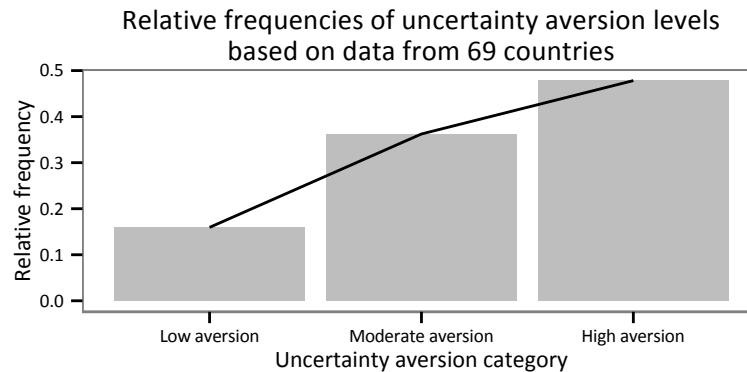
Figure 2 shows relative frequencies of the categories “low,” “moderate,” and “high” uncertainty aversion, on the basis of uncertainty aversion indices from 69 different countries. As can be clearly seen in Figure 2, the high-aversion group dominates,

nesses without ever changing the occupational status ‘self-employed,’ such that spells may overestimate the number of transitions.

⁶For the three questions, see page 190 of Hofstede et al. (2010).

⁷Respondent groups not consisting of IBM employees generate results strongly correlated to results obtained by questioning IBM employees (see Huang 2008).

FIGURE 2: Relative frequencies of low, moderate, and high uncertainty aversion



NOTE: Figure compiled on the basis of *k*-means clustering presented in Appendix C (see Table C.2). Data is taken from Hofstede (2015).

while the low-aversion group has the lowest relative frequency. Thus, I expect uncertainty aversion to be high while analyzing my model.

Research shows that uncertainty aversion has an impact on a number of variables. Without trying to provide a complete overview but only some examples: Huang (2007) shows that countries scoring higher on Hofstede's uncertainty aversion index, indicating higher aversion, export less to distant countries. Huang (2008) provides evidence that countries with higher uncertainty aversion grow slower in industrial sectors where information is scarce. Ramírez & Tadesse (2009) demonstrate that firms in countries with higher uncertainty aversion hold more cash. Inklaar & Yang (2012) show that financial crises have a stronger negative effect on investment in countries with higher uncertainty aversion.

I contribute to research on the impact of uncertainty aversion by adding a new effect to the literature: Uncertainty aversion may influence occupational choice decisions, and even more so, uncertainty aversion might explain phenomena like the earnings puzzle.

Summarizing strands of research relevant for the paper at hand, I find that most models in the literature on occupational choice operate with deterministic or risky rewards and do not aim to explain multiple transitions between occupations. By combining the strand of research on occupational choice and the strand of research on multi-armed bandits, I introduce a new model type with unknown stochastic rewards that is able to generate multiple transitions.

Contemporary occupational choice models do not provide a satisfactory explanation for the earnings puzzle. To obtain results consistent with the earnings puzzle, I draw on research on uncertainty-aversion effects, a third strand of research relevant for my paper. I contribute to the literature examining the earnings puzzle by showing that sufficiently high uncertainty aversion results in earnings-puzzle-like observations in an occupational choice model with partially unknown rewards. Empirical observations provide evidence for high uncertainty aversion and a low number of transitions between occupations. I use the joint evidence on uncertainty-aversion levels and transitions to assess the plausibility of model predictions.

3 The model

I consider an occupational choice model with a representative individual and partially unknown reward distributions. As this type of models can be most precisely expressed in the language of multi-armed bandit problems, I alternate between the terminology of occupational choice and bandit problems (gambling) throughout the paper. First, I introduce the model's setup. The model relies on a rather non-traditional outcome formulation with successes and nonsuccesses as rewards. Therefore, secondly, I discuss the type of rewards used in the model and relate it to conventional monetary rewards. Then, I examine the formation process of beliefs, thereby, finalizing the setup.

3.1 Setup

Think of a casino with a bandit. The bandit has three arms. Arm w represents wage work; arm e is entrepreneurship; and arm u is unemployment. Thus, there is a set of occupational options $\mathbb{O} := \{e, w, u\}$. Sometimes, it is more convenient to use a smaller set of options $\mathbb{O}_{-u} := \mathbb{O} \setminus u = \{e, w\}$.

A gambler can play the bandit in every time period from 1 until infinity, where time periods are indexed by $t \in \{0, 1, 2, \dots\}$. The three arms are independent. Only one arm can be selected in a particular period (an individual must not have more than one occupation at the same time). But, one arm must be pulled (an individual must have some occupation including unemployment). If the gambler selects an option $i \in \mathbb{O}$ in some period, I also say that the gambler samples from $i \in \mathbb{O}$. Switching from one arm to another does not result in any *direct* costs. If one arm is pulled, the remaining arms rest and cannot be observed (an individual can only gain information about an occupation by selecting it). Pulling an arm yields a reward: either a success, represented by $\pi = 1$, or a nonsuccess, represented by $\pi = 0$. Wage work generates an i.i.d. reward sequence $\{\Pi_{w,t}\}_{t=1}^{\infty}$ (π is a realization of Π) drawn from a Bernoulli distribution with probability mass function

$$f(\pi; \phi_w) = \phi_w^\pi (1 - \phi_w)^{1-\pi} \quad \text{for } \pi \in \{0, 1\}$$

where $\phi_w \in (0, 1)$ is the success probability of wage work. Similarly, entrepreneurship generates an i.i.d. reward sequence $\{\Pi_{e,t}\}_{t=1}^{\infty}$ drawn from a Bernoulli distribution with probability mass function $f(\pi; \phi_e)$ where $\phi_e \in (0, 1)$ is the success probability of entrepreneurship. Unemployment generates a nonsuccess with probability 1 in all periods.

Π_w , Π_e , and the unemployment reward $\Pi_{u,t} = 0$ cannot be observed in the same period. What can be observed is the result of using one particular arm (selecting a particular occupational option) in every period represented by the sequence $\{\Pi_t\}_{t=1}^{\infty}$. The gambler geometrically discounts rewards with a factor $\delta \in [0, 1)$, where $1 - \delta$ is the probability that at any given time the casino will close (equivalent to death or inability to work).

Parameters ϕ_w and ϕ_e are unknown. I assume that there is some prior knowledge

such that a prior on ϕ_w , respectively a prior on ϕ_e , can be specified. In fact, not knowing at least one parameter makes the decision problem more interesting since it does not allow for too simple decision strategies. For instance, in order to explore, i.e., to obtain new information, the gambler might be compelled to give up some allegedly secure income. Let $\mathbb{x}_{i,0}$ denote initial, imperfect, prior information about option $i \in \mathbb{O}_{-u}$. (Unemployment is a deterministic option such that information about it is always perfect.) Imperfect information is updated after every decision round. Consequently, there is information $\mathbb{x}_{i,t}$ about every option $i \in \mathbb{O}_{-u}$ in every period $t \in \{0, 1, \dots\}$. Let $\mathbb{X}_t := (\mathbb{x}_{e,t}, \mathbb{x}_{w,t})$ denote all information in a certain period.

Individuals are required to select a deterministic decision rule $\{d_t\}_{t=1}^{\infty}$ that specifies how to choose in every period. $d_t \in \mathbb{O}$ indicates the choice in period $t \in \{1, 2, \dots\}$. Rules continuously mixing are not allowed but mixing in a specific situation (for instance, a tie) is permitted. d_t can depend on all previous decisions $d_{t-1}, d_{t-2}, \dots, d_1$ and all past observations $\pi_{t-1}, \pi_{t-2}, \dots, \pi_1$ that are both captured by information \mathbb{X}_t , where the initialization \mathbb{X}_0 does not necessarily depend on actual decisions; and some additional parameter vector λ reflecting preferences, etc. It is assumed that the gambler follows an index rule. In every period $t \geq 0$, the gambler constructs an index $r_i(\mathbb{X}_t; \lambda) \in \mathbb{R}$ for every option $i \in \mathbb{O}$ and selects the option with the largest index in the next period. Let

$$\mathbb{O}_t^* := \{i \in \mathbb{O} : r_i(\mathbb{X}_t; \lambda) = \max\{r_e(\mathbb{X}_t; \lambda), r_w(\mathbb{X}_t; \lambda), r_u(\mathbb{X}_t; \lambda)\}\}$$

and

$$d_{t,r} := d_t(r_e(\mathbb{X}_{t-1}; \lambda), r_w(\mathbb{X}_{t-1}; \lambda), r_u(\mathbb{X}_{t-1}; \lambda))$$

Thus, we have

$$d_{t,r} = \begin{cases} \mathbb{O}_{t-1}^* & \text{if } |\mathbb{O}_{t-1}^*| = 1 \\ \text{each } i \in \mathbb{O}_{t-1}^* \text{ with probability } \frac{1}{|\mathbb{O}_{t-1}^*|} & \text{if } |\mathbb{O}_{t-1}^*| > 1 \end{cases} \quad (1)$$

Equation (1) can be further simplified by assuming that if $|\mathbb{O}_{t-1}^*| > 1$ and $u \in \mathbb{O}_{t-1}^*$, we should select $d_{t,r} = u$ —as unemployment is a no-effort option, while all other option require a certain effort.

Every decision $d_{1,r}, d_{2,r}, d_{3,r}, \dots$ will be associated with a reward $\Pi(d_{t,r}) \in \{0, 1\}$. The success history of an individual's career is given by

$$S_r := \sum_{t=1}^{\infty} \delta^{t-1} \Pi(d_{t,r}) \quad (2)$$

In this paper, I use the setup above to answer the following five questions:

1. What is the optimal strategy in a decision environment with elements of uncertainty, i.e., the decision rule maximizing the expected value of (2) conditional on prior information \mathbb{X}_0 , and what properties does the optimal decision rule have?

2. How do individuals constrained by uncertainty aversion decide and what properties does their behavior exhibit, especially in comparison to the optimal rule?
3. How does uncertainty aversion affect the propensity to select relatively less uncertain options? Here: The deterministic option u is of particular interest.
4. How does uncertainty aversion affect the correction of overconfidence with respect to entrepreneurship or, put differently, what is the impact of uncertainty aversion on learning efficiency?
5. How are transitions from one occupation to another affected by uncertainty aversion?

The answers to the five questions provide a benchmark in the form of optimal decisions; they provide intuition into negative effects of uncertainty aversion; they provide explanations for observations like the earnings puzzle; and they also generate a number of testable predictions.

3.2 Reward types

Instead of rewards obeying a Bernoulli distribution, we could also assume monetary rewards. Put differently, we could assume the following: Wage work generates an i.i.d. income sequence $\{V_{w,t}\}_{t=1}^{\infty}$ drawn from a distribution $\Psi_w(\boldsymbol{\theta}_w)$ where $V_{w,t} \in \mathbb{R}^+$. Entrepreneurship generates an i.i.d. income sequence $\{V_{e,t}\}_{t=1}^{\infty}$ drawn from a distribution $\Psi_e(\boldsymbol{\theta}_e)$ where $V_{e,t} \in \mathbb{R}$. Pulling the unemployment arm results in a fixed outcome, unemployment benefits or residual income, $v_u \in \mathbb{R}^+$. We cannot observe realizations of $V_{w,t}$ and $V_{e,t}$ at the same time but we can observe realizations of a sequence $\{V_t\}_{t=1}^{\infty}$ resulting from using a particular option in every period. Parameter vectors $\boldsymbol{\theta}_w$ and $\boldsymbol{\theta}_e$ are unknown but we can assume that there is some prior knowledge allowing the construction of prior distributions.

It is possible to link the formulation with monetary rewards to the Bernoulli-rewards formulation. Since success is a relative concept, we can use the secure unemployment benefit, or residual income, v_u , to distinguish successes from non-successes. Let $\pi^- \in \mathbb{R}^+$ denote a threshold value. Note that it is possible to define option-specific thresholds: $\pi_w^-, \pi_e^-, \pi_u^- \in \mathbb{R}^+$. A success can, for instance, be defined by

$$\pi_t = \begin{cases} 1 & \text{if } v_t - v_u \geq \pi^- \\ 0 & \text{if } v_t - v_u < \pi^- \end{cases}$$

where v_t is a realization of V_t ; and thresholds are assumed to be equal across options. This threshold definition of success simply states that if we scale down or understanding of success to a binary judgment, success in occupation $i \in \mathbb{O}$ is equivalent to the presence of an adequate monetary margin between residual income and the monetary income obtained in occupation $i \in \mathbb{O}$. Hence, the formulation with monetary rewards can be transformed into a formulation with success rewards, while the *nature of the problem* remains the same, as the gambler has to decide which occupation to select given that incomes are drawn from unknown distributions.

An additional aspect of the formulation with success rewards deserving attention is that it works with monetary and non-monetary incentives, and, thus, accounts for the importance of non-monetary rewards, because monetary and non-monetary incomes can be both transformed into a success variable. For instance, if we would like to model the effect of being your own boss, we could assume that the margin demanded for entrepreneurship is smaller than the margin demanded for wage work, i.e., $\pi_e^- < \pi_w^-$. This particular way to implement non-monetary aspects has the advantage that it does not introduce any conceptually new or *ad hoc* variables.

3.3 Formation of beliefs

I now model how beliefs about success probabilities are formed. Based on the formation process of beliefs, I introduce three measures to describe information about success probabilities: a measure of the amount of information denoted by $\gamma(\mathbf{x})$; the expected probability to succeed, $\mu(\mathbf{x})$; and the variance of information about the probability to succeed, $\sigma(\mathbf{x})^2$, which plays a prominent role in the definition of uncertainty given in the next section.

Let $\phi \in (0, 1)$ denote the unknown success probability of an arbitrary occupational option, excluding unemployment. Success Π obeys a Bernoulli distribution such that a success is returned with probability ϕ and a nonsuccess with probability $1 - \phi$; and such that $\mathbb{E}_{f(\pi; \phi)}[\Pi] = \phi$. We assume that ϕ is unknown but we let the gambler have a prior on ϕ . Furthermore, assume that the prior is beta. The beta distribution has the advantage that it is a proper conjugate prior in this setting. Hence, we have

$$\phi | \alpha, \beta \sim \mathfrak{B}(\alpha, \beta) \quad (3)$$

where \mathfrak{B} denotes the beta distribution, and $\alpha > 0$ and $\beta > 0$ are two hyperparameters. The beta distribution has density $B(\alpha, \beta)^{-1} \phi^{\alpha-1} (1 - \phi)^{\beta-1}$, where $B(\alpha, \beta) := \int_0^1 x^{\alpha-1} (1 - x)^{\beta-1} dx$ is the beta function.

How will prior beliefs change if the gambler observes a one-period success or nonsuccess? Let $P(\phi) = B(\alpha, \beta)^{-1} \phi^{\alpha-1} (1 - \phi)^{\beta-1}$ denote our prior belief and $P(\pi | \phi) = \phi^\pi (1 - \phi)^{1-\pi}$ denote the likelihood of observation $\pi \in \{0, 1\}$. According to Bayes' rule, the posterior is

$$P(\phi | \pi) = \frac{P(\pi | \phi)P(\phi)}{\int_0^1 P(\pi | \phi)P(\phi)d\phi} = \frac{\phi^{[\alpha+\pi]-1}(1-\phi)^{[\beta+1-\pi]-1}}{B(\alpha + \pi, \beta + 1 - \pi)}$$

corresponding to the density of a beta-distributed variable with parameters $\alpha + \pi$ and $\beta + 1 - \pi$. Hence, beliefs about success probability ϕ are updated according to

$$\mathfrak{B}(\alpha_t, \beta_t) \longrightarrow \begin{cases} \mathfrak{B}(\alpha_{t+1} = \alpha_t + \pi, \beta_{t+1} = \beta_t + 1 - \pi) & \text{if } \pi \text{ observed} \\ \mathfrak{B}(\alpha_{t+1} = \alpha_t, \beta_{t+1} = \beta_t) & \text{if } \pi \text{ not observed} \end{cases} \quad (4)$$

As the hyperparameters α and β fully describe the distribution of the success probability, we assume that $\mathbf{x}_{i,t} = (\alpha_{i,t}, \beta_{i,t})$ for all $i \in \mathbb{O}_{-u}$.

The hyperparameters and their updates are straightforward to interpret. Think of α_t and β_t as of two checklists evolving over time. α represents the number of successes and β the number of nonsuccesses. Given a success, α is increased by 1 and β stays constant. Given a nonsuccess, α stays constant and β is increased by 1. Initial parameters $\mathbf{x}_0 = (\alpha_0, \beta_0)$ are determined by pseudo observations (e.g., some historical data). A pseudo observation is, for instance, the success history of a peer who is considered to be similar with respect to some aspects (for instance, skills) important for occupational choice. However, pseudo observations are not necessarily generated by the same reward distribution as actual observations $\pi_{e,t}$ and $\pi_{w,t}$.

The checklist character of the hyperparameters allows the construction of simple measures. Let

$$\gamma(\mathbf{x}) := \alpha + \beta \quad (5)$$

Since every observed period either increases α or β , $\gamma(\mathbf{x})$ is a deterministic measure of the amount of information. The expected success probability is the number of successes divided by the number of observations:

$$\mu(\mathbf{x}) := \int_0^1 \phi f(\phi; \mathbf{x}) d\phi = \frac{\alpha}{\gamma(\mathbf{x})} \quad (6)$$

The variance of the information about the probability to succeed can be measured by

$$\sigma(\mathbf{x})^2 := \int_0^1 [\phi - \mu(\mathbf{x})]^2 f(\phi; \mathbf{x}) d\phi = \mu(\mathbf{x}) \frac{1 - \mu(\mathbf{x})}{1 + \gamma(\mathbf{x})} \quad (7)$$

An interesting point is how expectations and variance will develop given a success or nonsuccess. It is easy to see that a success strictly increases, $\mu(\alpha + 1, \beta) > \mu(\alpha, \beta)$, and a nonsuccess strictly decreases, $\mu(\alpha, \beta + 1) < \mu(\alpha, \beta)$, the expected success probability. The effect on variance depends on the ratio of successes to nonsuccesses. If the number of successes is much smaller than the number of nonsuccesses, the gambler will expect a nonsuccesses rather than a success.⁸ In such a situation, a success will increase the variance. A similar statement holds for the effect of a nonsuccess.

4 Rules for occupational choices under uncertainty

How will individuals solve their occupational choice problems given that outcomes are partially uncertain? In the first part of this section, I examine optimal decisions maximizing the expected sum of rewards given prior information. Then, I construct a decision rule accounting for uncertainty preferences in the second part of the section.

⁸More precisely, if $\alpha < ([9\beta^2 + 10\beta + 1]^{1/2} - \beta - 1)/4$, a success will increase the variance: $\sigma(\alpha + 1, \beta)^2 > \sigma(\alpha, \beta)^2$.

4.1 Decisions maximizing expected rewards

The most simple approach to evaluate the impact of preferences towards uncertainty is to compare uncertainty-preferences-affected decisions to optimal decisions. A prerequisite of such an approach is the existence of an optimal decision rule. A decision rule is said to be optimal if it induces decisions d_1, d_2, \dots maximizing the expected sum of discounted successes

$$\mathcal{S} := \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} \Pi(d_t) | \mathbb{x}_{e,0}, \mathbb{x}_{w,0} \right] \quad (8)$$

given some prior information $\mathbb{x}_{e,0} = (\alpha_{e,0}, \beta_{e,0})$ and $\mathbb{x}_{w,0} = (\alpha_{w,0}, \beta_{w,0})$. The best explanation of how to derive an optimal rule maximizing \mathcal{S} , which is simultaneously a proof of optimality, is presented by Weber (1992) and reproduced in Gittins et al. (2011).

PROPOSITION 1. (Optimal decision rule) *The optimal decision rule is an index rule selecting the option $i \in \mathbb{O}$ with the largest index $g_i(\mathbb{x}_t)$ in every period $t + 1$.⁹ Indices can be computed separately for each option such that $g_i(\mathbb{x}_t) = g(\mathbb{x}_{i,t})$. The index is the solution of the following optimal stopping problem:*

$$g(\mathbb{x}_{i,t}) = \sup_{t_s > t} \left\{ \frac{\mathbb{E} \left[\sum_{k=t+1}^{t_s} \delta^{k-1} \Pi_{i,k} | \mathbb{x}_{i,t} \right]}{\mathbb{E} \left[\sum_{k=t+1}^{t_s} \delta^{k-1} | \mathbb{x}_{i,t} \right]} \right\} \quad (9)$$

Proof. See Gittins et al. (2011) or Weber (1992). ■

Gittins & Jones (1974) were the first to solve the Bayesian multi-armed bandit problem with Bernoulli outcomes. Hence, we can refer to $g(\mathbb{x})$ as the Gittins index. Given an optimal stopping time t_s , the Gittins index in (9) is the ratio between the expected sum of discounted rewards and the expected sum of “discounted time.”

As unemployment is a known deterministic option, we get

$$g(\mathbb{x}_{u,t}) = \max_{t_s > t} \left\{ \frac{\sum_{k=t+1}^{t_s} \delta^{k-1} \Pi_{u,k}}{\sum_{k=t+1}^{t_s} \delta^{k-1}} \right\} = 0 \quad \text{for all } t$$

In the more general case of a stochastic option with an unknown success probability, the index does not have a known closed-form solution such that we have to rely on numerical approximations. Gittins et al. (2011) present different ways to approximate the index. A special case is myopia. If the probability to continue for one more period is $\delta = 0$, it is optimal to select the option with the largest expected success probability or $g(\mathbb{x}_{i,t}) = \mu_{i,t}$.¹⁰

⁹Note that rule g is not simply the optimal *index* rule but also the *best of all possible deterministic rules* (the best deterministic strategy) for the bandit model.

¹⁰The result stems from the fact that if $\delta = 0$, there will be no incentive to explore (the game ends in the next period). As only the incentive to exploit remains, it is optimal to select the currently most promising option (for a formal argument, see Gittins et al. 2011).

The Gittins index has three important properties:

PROPOSITION 2. (Properties of optimal rule) *The optimal rule, i.e., the decision rule induced by index g , has the following properties:*

(PG1) (Staying on a winner) *The Gittins index always stays on a winner, i.e., $g(\alpha + 1, \beta) > g(\alpha, \beta)$.*

(PG2) (Index boundaries) *The Gittins index has the following boundaries:*

$$\mu(\mathbf{x}) \leq g(\mathbf{x}) \leq \mu(\mathbf{x}) + \frac{\delta}{1 - \delta} \sigma(\mathbf{x})$$

(PG3) (Imperfect learning) *The optimal rule does not result in perfect learning: There is a nonzero probability that the best option is not selected for an infinite number of periods.*

Proof. For property (PG1), see Bellman (1956). For property (PG2) and (PG3), see Brezzi & Lai (2000). ■

According to property (PG2), a Gittins-index strategy will never select unemployment, as $g(\mathbf{x}_{u,t}) = 0$ for all t but $g(\mathbf{x}_{i,t}) \geq \mu(\mathbf{x}_{i,t}) > 0$ for all $i \in \mathbb{O}_{-u}$ and t . Property (PG3), the fact that optimal learning and perfect information are not equivalent, has one important implication: Counterfactuals constructed under the assumption of perfect information are not an appropriate benchmark. In a realistic setting, decision makers first need to learn what the counterfactual is. Since optimal learning does not result in the identification of the best option with a probability of 100%, even optimally behaving individuals might not decide on the basis of the best counterfactual if success probabilities are unknown.

4.2 Decisions affected by uncertainty preferences

The optimal rule does not take into account that individuals might be uncertainty-averse or for preferences towards uncertainty in general. Without doubt, there are numerous ways to formalize uncertainty and uncertainty preferences.¹¹ Additionally, uncertainty-preferences-affected behavior can manifest in different facets. The choice of definitions below is guided by considerations of simplicity. Therefore, uncertainty and preferences are defined in a rather general way. With respect to preferences towards uncertainty, I only require consistency with the following empirical result:

WEAK ELLSBERG REQUIREMENT. (WER) *Individuals tend to prefer an option with a known success probability to an uncertain option.*

¹¹The reason why I do not use traditional concepts as, for example, subjective expected utility (Savage 1954) is that they do not differentiate between risk and uncertainty. However, the difference is substantial since in case of risk reward distributions are known and learning about reward distributions is pointless, whereas in case of uncertainty learning improves decisions.

WER is based on Ellsberg's (1961) findings and is general enough to capture the idea of uncertainty aversion but does not impose a particular way to model it. First, I construct a rule accounting for uncertainty preferences and establish some of its properties. For a lack of a better term, I refer to the rule as the 'behavioral rule,' to contrast it from the optimal rule, or simply as 'rule b.' Then, I discuss empirical validity of the rule by relating it to experimental evidence.

The first step in the construction of a rule accounting for uncertainty preferences is the following definition of uncertainty¹²:

DEFINITION 1. (Uncertainty) The uncertainty of an option $i \in \mathbb{O}$ is reflected by the standard deviation, $\sigma(\mathbf{x}_{i,t})$, of the prior distribution on success probability ϕ_i . If $\sigma(\mathbf{x}_{i,t}) > \sigma(\mathbf{x}_{j,t})$, we say that, in period $t \in \{0, 1, 2, \dots\}$, option i is more uncertain than option j (and *vice versa*), where $i \neq j$. If an option j is not uncertain, we have $\sigma(\mathbf{x}_{j,t}) = 0$ for all t .

Using the standard deviation $\sigma(\mathbf{x})$ to measure uncertainty is straightforward as, according to Chebyshev's inequality, the probability of an estimation error, i.e., $|\phi_i - \mu(\mathbf{x}_{i,t})|$ is larger than or equal to some $\phi_\Delta > 0$, is bounded by

$$\mathbb{P} \{ |\phi_i - \mu(\mathbf{x}_{i,t})| \geq \phi_\Delta \} \leq \frac{\sigma(\mathbf{x}_{i,t})}{\phi_\Delta} \quad \text{for } \phi_\Delta > \sigma(\mathbf{x}_{i,t})$$

where the boundary, $\sigma(\mathbf{x}_{i,t})/\phi_\Delta$, decreases if uncertainty decreases.

It is assumed that the behavioral rule is an index rule where the option with the largest index $b(\mathbf{x}_{i,t}; \lambda, \gamma_b)$ is selected in every period $t + 1$. In contrast to the optimal rule (the Gittins index), the behavioral rule does not only depend on observations $\mathbf{x}_{i,t} = (\alpha_{i,t}, \beta_{i,t})$ but also on a parameter $\lambda \in \Lambda$ representing preferences towards uncertainty and on a normalizing constant γ_b representing an uninformative state of knowledge. In our beta-Bernoulli setting, the state with no information can be represented by $\gamma_b = \alpha_0 + \beta_0 = 2$, as this implies a uniform prior distribution, $\mathfrak{B}(1, 1)$, for the probability to succeed, ϕ .

λ is a psychological primitive. It can reflect some measure of uncertainty tolerance (Sherman 1974) or a measure of uncertainty avoiding behavior like the uncertainty avoidance index developed by Hofstede et al. (2010). λ might also be associated with the Openness factor in the Five Factor Model of personality (see McCrae & John 1992; John & Srivastava 2001).

I assume that uncertainty preferences do not directly affect expected probabilities to succeed and that the effect of uncertainty preferences, which is denoted by $\psi(\cdot)$, is additive such that the behavioral index can be written as

$$b(\mathbf{x}_{i,t}; \lambda, \gamma_b) := \mu(\mathbf{x}_{i,t}) + \psi(\mathbf{x}_{i,t}; \lambda, \gamma_b) \quad (10)$$

A sensible requirement is that we have $\psi(\mathbf{x}_{i,t}; \lambda, \gamma_b) = 0$ if an option is not uncertain.¹³ Furthermore, it is assumed that if an option is uncertain, the following

¹²The definition is inspired by conceptually similar definitions of risk (see, for instance, Tobin 1958).

¹³As an infinite amount of information is equivalent to perfect information, an additional sensible

relations hold:

$$\psi(\mathbb{x}_{i,t}; \lambda, \gamma_b) < 0 \iff b(\mathbb{x}_{i,t}) < \mu(\mathbb{x}_{i,t}) \text{ if } \begin{array}{l} \psi(\mathbb{x}_{i,t}; \lambda, \gamma_b) \text{ represents} \\ \text{uncertainty aversion} \end{array} \quad (11)$$

$$\psi(\mathbb{x}_{i,t}; \lambda, \gamma_b) = 0 \iff b(\mathbb{x}_{i,t}) = \mu(\mathbb{x}_{i,t}) \text{ if } \begin{array}{l} \psi(\mathbb{x}_{i,t}; \lambda, \gamma_b) \text{ represents} \\ \text{uncertainty neutrality} \end{array} \quad (12)$$

$$\psi(\mathbb{x}_{i,t}; \lambda, \gamma_b) > 0 \iff b(\mathbb{x}_{i,t}) > \mu(\mathbb{x}_{i,t}) \text{ if } \begin{array}{l} \psi(\mathbb{x}_{i,t}; \lambda, \gamma_b) \text{ represents} \\ \text{uncertainty affinity} \end{array} \quad (13)$$

where arguments of $b(\mathbb{x}_{i,t}; \lambda, \gamma_b)$ not depending on time are suppressed for notational convenience.

Assumptions underlying (11)–(13) are consistent with a rather simple and general definition of uncertainty preferences:

DEFINITION 2. (Uncertainty preferences) There are three types of attitudes towards uncertainty (preferences are transitive):

- | | |
|------------|--|
| Neutrality | Given options with the same expected success probability but different uncertainty levels, an uncertainty-neutral gambler is always indifferent between all options. |
| Affinity | Given two options 1 and 2 where both options have the same expected success probability but there is more information about option 1 than 2 and option 2 is more uncertain than 1, an uncertainty-affine gambler will always strictly prefer option 2 to option 1. |
| Aversion | Given two options 1 and 2 where both options have the same expected success probability but there is more information about option 1 than 2 and option 2 is more uncertain than 1, an uncertainty-averse gambler will always strictly prefer option 1 to option 2. |

Without loss of generality, let us assume that $\Lambda = \mathbb{R}$. Uncertainty neutrality is represented by $\lambda = 0$. Uncertainty affinity is represented by $\lambda \in \mathbb{R}^+$. Finally, uncertainty aversion is represented by $\lambda \in \mathbb{R}^-$. Uncertainty neutrality can be used as a reference point to construct absolute uncertainty aversion: An individual 1 is assumed to be more uncertainty averse than an individual 2 if $\lambda_1 < \lambda_2$ where $\lambda_1, \lambda_2 \in \mathbb{R}^-$ are uncertainty preferences associated with individual 1 and 2. For the remainder of the paper, in line with findings presented in Section 2, I assume that uncertainty aversion is the most realistic preference towards uncertainty.

For further analysis, it will be convenient to have an explicit definition of the uncertainty preferences effect. I propose the following function:

DEFINITION 3. (Uncertainty preferences effect) Uncertainty preferences are represented by

$$\psi(\mathbb{x}_{i,t}; \lambda, \gamma_b) := \lambda \frac{\gamma_b}{\gamma(\mathbb{x}_{i,t})} \sigma(\mathbb{x}_{i,t}) \quad (14)$$

requirement is $\mathbb{P} \{ \lim_{h_i \rightarrow \infty} \psi(\mathbb{x}_{i,h_i}; \lambda, \gamma_b) = 0 \} = 1$, where h_i denotes the number of times option $i \in \mathbb{O}$ is selected.

Definition 3 has a number of advantages. First, it is very simple—in effect, it is just a dynamic weighting of uncertainty. Second, it is consistent with experimental findings, discussed further below. Thirdly, the partial weight $\gamma_b \gamma(x_{i,t})^{-1}$ captures habituation effects: Camerer & Ho (1999) establish a high explanatory performance of learning algorithms featuring self-reinforcement. In a situation where the gambler is faced with two options with the same expected success probability and same uncertainty, it is reasonable to assume that the gambler will still treat the two options differently if the amount of information is different across options, which is ensured by (14), as the behavioral rule with function (14) assumes that uncertainty-averse individuals will exhibit a tendency to prefer the option with more information. Fourthly and lastly, as, according to the behavioral rule with uncertainty effect (14) and given a relatively high level of uncertainty aversion, there will be a tendency to select an option with less uncertainty over an option with more uncertainty even when the option with more uncertainty has a superior expected probability to succeed, the particular choice of (14) is consistent with the weak Ellsberg requirement (WER). It is also obvious that a behavioral rule with uncertainty effect (14) is consistent with Definition 2.¹⁴

The behavioral rule has four important properties.

PROPOSITION 3. (Properties of behavioral rule) *The behavioral rule, i.e., rule b , has the following properties:*

- (PB1) (Absorption) *If at some point in time it is weakly preferred to select unemployment, unemployment will be selected in all consecutive periods.*
- (PB2) (Unusual types) *There exist two unusual types of individuals. Type-G individuals can sometimes gamble for the sake of gambling by preferring an uncertain option to a sure success. Type-P individuals can prefer to be unemployed since this choice yields an uncertainty-free income such that they are paralyzed in their choices by uncertainty.*
- (PB3) (Not staying on a winner) *If a success was not expected due to a large number of previous nonsuccesses, the gambler might not stay on a winner.*
- (PB4) (Imperfect learning) *Even when all prior observations $x_{i,0}$ for all $i \in \mathbb{O}_{-u}$ are sampled from a Bernoulli distribution with a correct success probability ϕ_i , the behavioral rule will with a nonzero probability never identify the best option.*

Proof. See Appendix A. ■

Properties of the behavioral and the optimal rule are quite similar (e.g., both rules do not lead to perfect learning) but the properties of the behavioral rule might be considered as slightly more realistic. For instance, it does not make sense to stay on a winner (the optimal rule always stays on a winner) if a win is confusing because unexpected. We also might expect the existence of unusual individual

¹⁴Consider two options 1 and 2 where $\mu_1 = \mu_2$, $\gamma_1 > \gamma_2$, and $\sigma_1 < \sigma_2$. If $\lambda = 0$, we have $b(x_1) = b(x_2)$. If $\lambda \in \mathbb{R}^+$, we get $b(x_1) < b(x_2)$. If $\lambda \in \mathbb{R}^-$, we have $b(x_1) > b(x_2)$.

types (e.g., gamblers who gamble for the sake of gambling), which is allowed by the behavioral rule. Careers are path-dependent (see, e.g., Bernhardt, Morris, Handcock & Scott 2001) and the behavioral rule also shows many features of path dependency. For instance: There is an absorbing state (unemployment¹⁵); decisions are self-reinforcing; due to chance, the best option might never be identified.

Experimental evidence on Bernoulli bandits is rather scarce but there are two particularly interesting findings provided by Steyvers, Lee & Wagenmakers (2009) and Gans, Knox & Croson (2007). First, Steyvers et al. (2009) calculate different behavioral characteristics of individuals playing with Bernoulli bandits. A characteristic is, for instance, the number of times individuals selected arms with fewer success and nonsuccesses than alternative arms, which can be interpreted as exploratory behavior. Steyvers et al. (2009) could not establish significant correlations between characteristics and personality. In Appendix B, I successfully replicate this finding in a simulated experiment using the setting of Steyvers et al. (2009) and the behavioral rule. Interestingly enough, the result only holds if the share of uncertainty-averse individuals is sufficiently large in the group under examination.

Second, Gans et al. (2007) apply different decision rules to data generated by the decisions of individuals who played with Bernoulli bandits. They establish a ranking of the explanatory performance of different rules. The ranking is successfully replicated in Appendix B using decision data generated by the behavioral rule. It is only possible to replicate the ranking if uncertainty aversion is pronounced and individuals mostly follow the behavior rule without making too many mistakes. Thus, the behavioral rule with uncertainty aversion is supported by available experimental evidence on Bernoulli bandits.

In summary: I consider two decision rules. The first rule is individually optimal, whereas the second rule allows for the influence of preferences towards uncertainty. The latter rule is more realistic than the optimal rule. However, the optimal rule is a proper benchmark for decision making and, especially, the efficiency of learning.

5 The impact of uncertainty aversion on career decisions and overconfidence

Having introduced the model's setup, an individually optimal decision rule, and a decision rule accounting for uncertainty preferences, I now examine different aspects of occupational choices, in this way, presenting the main theoretical findings of the paper. There are four interesting decision aspects deserving a closer examination.

First, as unemployment has special features, compared to wage work and entrepreneurship that are neither certain nor absorbing, the propensity to select unemployment is of interest. Second, as entrepreneurs are usually overconfident, it

¹⁵Verbruggen, van Emmerik, Van Gils, Meng & de Grip (2015) find that underemployment has negative long-term effects on career success. A similar effect is present in my model, as an individual selecting unemployment effectively limits her future set of options.

is interesting how uncertainty aversion impacts the ability to correct false beliefs. Third, the main assumption of the paper is that individuals transition between occupations, i.e., occupational choices are not final decisions. Thus, an interesting aspect is the link between transitions and uncertainty preferences. Fourth, an aspect of interest are sampling frequencies: As uncertainty aversion may result in over- or undersampling from some options and sampling frequencies are observable, the analysis of sampling produces testable predictions.

This section consists of three parts. Propensities to select unemployment and overconfidence can be discussed in the model laid out in Section 3. I refer to this version as the 'full model.' Thus, in the first part of the section, I use the full model to derive results on the propensity to select unemployment and the correction of false beliefs. However, transitions are hard to discuss in the full model because the potential number of transitions is infinite. In such a case, a simple approach is to modify the full model to make it suitable to discuss a certain aspect. In the second part of the section, I examine transitions and sampling by adding the assumption of predetermined absorption, restricting the maximal number of transitions, to the full model. The result is a second version of the bandit model.

A counter argument to such an approach is that a modification might *overemphasize* some aspects that would not play a large role without the modification—for instance, uncertainty aversion may have a large impact on transitions in the modified model but a negligible impact in the full model. For this reason, I simulate a version of the full model in the third part of the section, resulting in a third model version, to verify that results derived in the model with predetermined absorption also apply to the full model. In other words, the third part of the section brings together insights from the first and second model version and builds up intuition by illustrating how derived propositions manifest themselves in a simulated version of the full model.

5.1 Choices in full model

With respect to the propensity to select unemployment, I obtain the following proposition:

PROPOSITION 4. (Selecting sure options) *Assume that $\mathbb{x}_{e,0}$ and $\mathbb{x}_{w,0}$ are generated by a fixed number of draws $\gamma_{e,0}$ and $\gamma_{w,0}$ from two independent Bernoulli distributions. Let $\mathbb{s}_r := \mathbb{P}\{d_{1,r} = u\}$ denote the propensity to select the sure option unemployment in the first period and let $\mathbb{s}_{r,T} := \mathbb{P}\{d_{t,r} = u \text{ for some } t \in \{1, 2, \dots, T\}\}$ denote the propensity to select the sure option unemployment in some period $t \in \{1, 2, \dots, T\}$. The behavioral rule has a (weakly) higher propensity to select a sure option in the first period:*

$$\mathbb{s}_b \geq \mathbb{s}_g$$

where the inequality is strict if λ is sufficiently small. The behavioral rule has a (weakly) higher propensity to select a sure option in any period $t \in \{1, 2, \dots, T\}$:

$$\mathbb{s}_{b,T} \geq \mathbb{s}_{g,T}$$

where the inequality is strict if λ is sufficiently small. Furthermore, given two different levels of uncertainty aversion λ' , $\lambda < 0$ where $\lambda' < \lambda$, we have

$$\mathbb{s}_b(\lambda') \geq \mathbb{s}_b(\lambda), \quad \mathbb{s}_{b,T}(\lambda') \geq \mathbb{s}_{b,T}(\lambda)$$

where inequalities are strict if λ is sufficiently large and λ' is sufficiently smaller than λ .

Proof. Let $\mu_{i,t} := \mu(\mathbb{x}_{i,t})$, $\sigma_{i,t} := \sigma(\mathbb{x}_{i,t})$, and $r_{i,t} := r(\mathbb{x}_{i,t}; \dots)$, where $r(\mathbb{x}_{i,t}; \dots)$ is either the behavioral or the optimal index. By independence of options, we have $\mathbb{P}\{d_{t,g} = u\} = \mathbb{P}\{g_{e,t-1} \leq 0\} \mathbb{P}\{g_{w,t-1} \leq 0\} = [1 - \mathbb{P}\{g_{e,t-1} > 0\}][1 - \mathbb{P}\{g_{w,t-1} > 0\}]$. Using the boundaries from (PG2), we obtain $\mathbb{P}\{g_{e,t} > 0\} = \mathbb{P}\{g_{w,t} > 0\} = 1$ for all t such that $\mathbb{P}\{d_{t,g} = u\} = 0$ for all $t > 0$. Consequently, we have $\mathbb{s}_g = \mathbb{s}_{g,T} = 0$ for all T . To compute \mathbb{s}_b , note that

$$\mathbb{P}\{d_{1,b} = u\} = \mathbb{P}\{\mu_{e,0} + \lambda \gamma_b \sigma_{e,0} \gamma_{e,0}^{-1} \leq 0\} \mathbb{P}\{\mu_{w,0} + \lambda \gamma_b \sigma_{w,0} \gamma_{w,0}^{-1} \leq 0\}$$

Let $\mathbb{x}_{e,0} = (X_e, \gamma_{e,0} - X_e)$ and $\mathbb{x}_{w,0} = (X_w, \gamma_{w,0} - X_w)$ where X_e and X_w denote the sum of (prior) successes in entrepreneurship and wage work given $\gamma_{e,0}$ and $\gamma_{w,0}$ observations. Let $Y_e := X_e [\gamma_b \sigma(X_e, \gamma_{e,0} - X_e)]^{-1}$ and $Y_w := X_w [\gamma_b \sigma(X_w, \gamma_{w,0} - X_w)]^{-1}$. Using (6), we obtain

$$\mathbb{P}\{d_{1,b} = u\} = \mathbb{P}\{Y_e \leq -\lambda\} \mathbb{P}\{Y_w \leq -\lambda\} = F_{Y_e}(-\lambda) F_{Y_w}(-\lambda)$$

Now notice that $\sigma(X, \gamma_0 - X)$ takes a maximum at $\mathbb{x} = (\gamma_0/2, \gamma_0/2)$ and a minimum at $\mathbb{x} = (0, \gamma_0)$ and $\mathbb{x} = (\gamma_0, 0)$, where the minimum is zero. Hence:

$$\frac{1}{\gamma_b \sigma(\gamma_{i,0}/2, \gamma_{i,0}/2)} \leq Y_i \leq \frac{\gamma_{i,0} - 1}{\gamma_b \sigma(\gamma_{i,0} - 1, 1)} \quad (15)$$

Consequently, we can always find a

$$\lambda < \lambda_{\mathbb{s}}^- := -\max \left\{ \frac{1}{\gamma_b \sigma(\gamma_{e,0}/2, \gamma_{e,0}/2)}, \frac{1}{\gamma_b \sigma(\gamma_{w,0}/2, \gamma_{w,0}/2)} \right\}$$

such that we have $F_{Y_e}(-\lambda), F_{Y_w}(-\lambda) > 0$ and, thus, $\mathbb{s}_b > \mathbb{s}_g$. Furthermore, if

$$\lambda < \lambda_{\mathbb{s}}^+ := -\max \left\{ \frac{\gamma_{e,0} - 1}{\gamma_b \sigma(\gamma_{e,0} - 1, 1)}, \frac{\gamma_{w,0} - 1}{\gamma_b \sigma(\gamma_{w,0} - 1, 1)} \right\}$$

we obtain $F_{Y_e}(-\lambda) = F_{Y_w}(-\lambda) = 1$ such that $\mathbb{s}_b = 1$. Let $\lambda', \lambda \in \mathbb{R}^-$. As $F_{Y_e}(-\lambda)$ and $F_{Y_w}(-\lambda)$ are both nondecreasing, we can always find a $\lambda' < \lambda$ such that $\mathbb{s}_b(\lambda') \geq \mathbb{s}_b(\lambda)$, where the inequality is strict if λ' is sufficiently smaller than λ . To assess $\mathbb{s}_{b,T}$, note that $1 - \mathbb{s}_{b,T}$ is the probability that unemployment is never selected in $1, \dots, T$. Let $\tau_{e,t} \geq 0$ denote the number of times the behavioral rule samples from entrepreneurship and $\tau_{w,t} \geq 0$ the number of times the rule samples from wage work up to period t . Notice that $\tau_{e,t} + \tau_{w,t} \leq T - 1$. Let $\mathbb{x}_{e,t-1} = (X_e + X_{e,t-1}^*, \gamma_{e,0} + \tau_{e,t-1} -$

$X_e - X_{e,t-1}^*$) and $\mathbb{x}_{w,t-1} = (X_w + X_{w,t-1}^*, \gamma_{w,0} + \tau_{w,t-1} - X_w - X_{w,t-1}^*)$, where $X_{e,t}$ and $X_{w,t}$ is the sum of successes in entrepreneurship and wage work given $\tau_{e,t}$ and $\tau_{w,t}$ observations. Define $Y_{e,t-1}^* := (X_e + X_{e,t-1}^*) \left[\gamma_b \sigma(X_e + X_{e,t-1}^*, \gamma_{e,0} + \tau_{e,t-1} - X_e - X_{e,t-1}^*) \right]^{-1}$ and $Y_{w,t-1}^* := (X_w + X_{w,t-1}^*) \left[\gamma_b \sigma(X_w + X_{w,t-1}^*, \gamma_{w,0} + \tau_{w,t-1} - X_w - X_{w,t-1}^*) \right]^{-1}$. In analogy to (15), we get

$$\lambda_{s,t}^- \leq Y_{i,t-1}^* \leq \lambda_{s,t}^+ \quad (16)$$

$$\lambda_{s,t}^- := \frac{1}{\gamma_b \sigma([\gamma_{i,0} + \tau_{i,t-1}]/2, [\gamma_{i,0} + \tau_{i,t-1}]/2)}, \quad \lambda_{s,t}^+ := \frac{\gamma_{i,0} + \tau_{i,t-1} - 1}{\gamma_b \sigma(\gamma_{i,0} + \tau_{i,t-1} - 1, 1)}$$

As $\sigma([\gamma_{i,0} + \tau_{i,t-1}]/2, [\gamma_{i,0} + \tau_{i,t-1}]/2)$ decreases in $\tau_{i,t-1}$, we must have $\lambda_{s,1}^- \geq \lambda_{s,2}^- \geq \dots \geq \lambda_{s,T}^-$. As $\gamma_{i,0} + \tau_{i,t-1} - 1$ increases in $\tau_{i,t-1}$ and $\sigma(\gamma_{i,0} + \tau_{i,t-1} - 1, 1)$ decreases in $\tau_{i,t-1}$, we must have $\lambda_{s,1}^+ \leq \lambda_{s,2}^+ \leq \dots \leq \lambda_{s,T}^+$. Furthermore, note that $\tau_{i,t-1} \leq t - 1$, as an option cannot be observed for more than $t - 1$ periods. Hence:

$$\lambda_{s,t}^- \geq \hat{\lambda}_{s,t}^- := \frac{1}{\gamma_b \sigma([\gamma_{i,0} + t]/2, [\gamma_{i,0} + t]/2)}, \quad \lambda_{s,t}^+ \leq \hat{\lambda}_{s,t}^+ := \frac{\gamma_{i,0} + t - 1}{\gamma_b \sigma(\gamma_{i,0} + t - 1, 1)}$$

where $\hat{\lambda}_{s,1}^- \geq \hat{\lambda}_{s,2}^- \geq \dots \geq \hat{\lambda}_{s,T}^-$ and $\hat{\lambda}_{s,1}^+ \leq \hat{\lambda}_{s,2}^+ \leq \dots \leq \hat{\lambda}_{s,T}^+$. Thus, if

$$\lambda > \lambda_{-,T} := - \min \left\{ \frac{1}{\gamma_b \sigma([\gamma_{e,0} + T]/2, [\gamma_{e,0} + T]/2)}, \dots \right. \\ \left. \dots, \frac{1}{\gamma_b \sigma([\gamma_{w,0} + T]/2, [\gamma_{w,0} + T]/2)} \right\}$$

where $\lambda_{-,T} < 0$, unemployment will never be selected in any period $t \in \{1, \dots, T\}$ and we obtain $\mathbb{s}_{b,T} = 0$. However, if

$$\lambda < \lambda_{+,T} := - \max \left\{ \frac{\gamma_{e,0} + T - 1}{\gamma_b \sigma(\gamma_{e,0} + T - 1, 1)}, \frac{\gamma_{w,0} + T - 1}{\gamma_b \sigma(\gamma_{w,0} + T - 1, 1)} \right\}$$

where $\lambda_{+,T} < \lambda_{-,T}$, unemployment will be selected with probability 1 such that $\mathbb{s}_{b,T} = 1$. Hence, there must exist a $\lambda < 0$ such that $\mathbb{s}_{b,T} > 0$, respectively $\mathbb{s}_{b,T} > \mathbb{s}_{g,T}$. Furthermore, given that $\lambda', \lambda \in \mathbb{R}^-$, there must exist a $\lambda' < \lambda$ such that $\mathbb{s}_{b,T}(\lambda') \geq \mathbb{s}_{b,T}(\lambda)$, where the inequality is strict if λ' is sufficiently smaller than λ . ■

Summarizing Proposition 4, it can be said that if the level of uncertainty aversion is sufficiently high, uncertainty-averse individuals will select unemployment with a higher than optimal probability, where the optimal probability is zero.

We know that entrepreneurs differ from other individuals (Zhao & Seibert 2006; Rauch & Frese 2007). A particular characteristic of entrepreneurs is overconfidence (Busenitz & Barney 1997; Camerer & Lovallo 1999; Bernardo & Welch 2001; Koellinger et al. 2007). I define overconfidence related to entrepreneurship as follows.

DEFINITION 4. (Overconfidence related to entrepreneurship) An individual is said to be overconfident with respect to entrepreneurship if she assigns a higher expected success probability to entrepreneurship than to wage work, $\mu(\mathbb{x}_{e,0}) > \mu(\mathbb{x}_{w,0})$, and

is less uncertain about success in entrepreneurship than in wage work, $\sigma(\mathbb{x}_{e,0}) < \sigma(\mathbb{x}_{w,0})$, but entrepreneurship generates a success with a lower probability than expected, $\phi_e < \mu(\mathbb{x}_{e,0})$.

Griffin & Varey (1996) argue that overconfidence expresses itself in two different ways:

[Type-1 overconfidence] One type, the most dramatic, is optimistic overconfidence, the tendency to overestimate the likelihood that one's favored outcome will occur. (Griffin & Varey 1996, p. 228)

[Type-2 overconfidence] A second type of overconfidence is the overestimation of one's knowledge (more generally, the overestimation of the validity of one's judgment) (Griffin & Varey 1996, p. 228)

Definition 4 covers both types of overconfidence. Type-1 overconfidence, the tendency to overestimate success probabilities, is reflected by the overestimation of the success probability of entrepreneurship, $\mu(\mathbb{x}_{e,0}) > \phi_e$. Type-2 overconfidence, the overestimation of one's knowledge, is captured by a smaller uncertainty about success probabilities in entrepreneurship than in wage work, $\sigma(\mathbb{x}_{e,0}) < \sigma(\mathbb{x}_{w,0})$.

Overconfidence does not automatically lead to an observation like the earnings puzzle, as overestimating the success probability of entrepreneurship does not rule out that entrepreneurship is the best option. Therefore, additionally assume that $\phi_w > \phi_e$ such that there is overconfidence with respect to entrepreneurship, while wage work has a higher success probability than entrepreneurship. Given this constellation, beliefs about the superiority of entrepreneurship are objectively false.

Assume that, driven by overconfidence with respect to entrepreneurship, the individual decides for entrepreneurship in the first period. After the individual selects entrepreneurship, she observes a signal in the form of success observations $\pi_{e,1}, \pi_{e,2}, \dots$, where $\pi_{e,t}$ is drawn from a Bernoulli distribution with success probability ϕ_e . To make the signal unambiguous, assume that it consists only of nonsuccesses (such a signal is consistent with $\phi_w > \phi_e$). The correction of overconfidence requires a transition to wage work.¹⁶ Proposition 5 depicts how transitions leading to a potential correction of overconfidence and uncertainty aversion are related:

PROPOSITION 5. (Overconfidence and uncertainty aversion) *Assume that an individual who is overconfident with respect to entrepreneurship receives a clear signal, $\mathbf{s}_{e,t}^* = (\pi_{e,1} = 0, \pi_{e,2} = 0, \dots, \pi_{e,t} = 0)$, that she overestimated the success probability of entrepreneurship and that wage work might be superior to entrepreneurship. Furthermore, assume that the following four conditions hold:*

- (C1) $\mu(\mathbb{x}_{e,0}) > \mu(\mathbb{x}_{w,0}) + \delta(1 - \delta)^{-1}\sigma(\mathbb{x}_{w,0});$
- (C2) $\sigma_e < \sigma(\mathbb{x}_{w,t})$ for all t and $\sigma_e < (1 - \delta)\delta^{-1}\mu(\mathbb{x}_{w,0})$, where $\sigma_e := \max_{k \geq 0} \sigma(\alpha_{e,0}, \beta_{e,0} + k);$
- (C3) $\lambda > -\alpha_{w,0}[\gamma_b \sigma(\mathbb{x}_{w,0})]^{-1};$

¹⁶Though, one transition might not be enough because the individual might return to the inferior option entrepreneurship. However, false beliefs cannot be corrected without a transition at all.

$$(C4) \quad \alpha_{e,0} > [\gamma_{e,0}/\gamma_{w,0}]\alpha_{w,0} \text{ and } \gamma_{e,0} > \gamma_{w,0}.$$

Let \mathbf{t}_g denote the period where an individual following a Gittins-index strategy experiences the incentive to abandon entrepreneurship and transition to wage work for the first time. Let \mathbf{t}_b denote the equivalent period for an individual following a behavioral-index strategy. If uncertainty aversion is sufficiently high, we have

$$\mathbf{t}_b > \mathbf{t}_g$$

such that $\mathbf{s}_{e,t}^*$ signaling inferiority of entrepreneurship, compared to wage work, needs longer than optimal to induce a transition to wage work. Furthermore, given two levels of uncertainty aversion λ' , $\lambda < 0$, an individual with higher uncertainty aversion will need longer to transition to entrepreneurship such that

$$\mathbf{t}_b(\lambda') > \mathbf{t}_b(\lambda)$$

if $\lambda' < \lambda$.

Proof. Let $\mu_{i,t} := \mu(\mathbb{x}_{i,t})$, $\sigma_{i,t} := \sigma(\mathbb{x}_{i,t})$, and $r_{i,t} := r(\mathbb{x}_{i,t}; \dots)$ [$r(\mathbb{x}_{i,t}; \dots)$ is either the behavioral or the optimal index]. Using (C1) and the boundaries in (PG2), we can establish $g_{e,0} > g_{w,0}$ such that $d_{1,g} = e$. By overconfidence, as given by Definition 5, $b_{e,0} > b_{w,0}$ and, thus, $d_{1,b} = e$. Given (C3), $b_{w,0} > 0$. Note that, given signal $\mathbf{s}_{e,t}^*$, $\sigma_e = \sigma_{e,t_\sigma}$, where $t_\sigma = -\beta_{e,0} + [9\alpha_{e,0}^2 + 10\alpha_{e,0} + 1]^{1/2}/4 - \alpha_{e,0}/4 - 1/4$. Consequently, $\mu_{e,t} + \delta(1-\delta)^{-1}\sigma_{e,t} \leq \mu_{e,t} + \delta(1-\delta)^{-1}\sigma_e$. To establish $\mathbf{t}_b > \mathbf{t}_g$, I show that $\mathbf{t}_g \in (0, \mathbf{t}_g^+]$, $\mathbf{t}_b \in [\mathbf{t}_b^-, \infty)$, and $\mathbf{t}_b^- > \mathbf{t}_g^+$ given some $\lambda < 0$ consistent with (C3).¹⁷ Using the boundaries in (PG2), we can establish that a transition incentive will definitely occur for the first time when $\mu_{e,t} + \delta(1-\delta)^{-1}\sigma_e = \mu_{w,0}$ but not later. Given signal $\mathbf{s}_{e,t}^*$, we obtain $\mu_{e,t} = \alpha_{e,0}(\gamma_{e,0} + t)^{-1}$. Hence, we have $\mathbf{t}_g \leq \mathbf{t}_g^+$, where

$$\mathbf{t}_g^+ = \frac{(1-\delta)\alpha_{e,0} - \mathbf{a}_0\gamma_{e,0}}{\mathbf{a}_0}, \quad \mathbf{a}_0 := (1-\delta)\mu_{w,0} - \delta\sigma_e \quad (17)$$

Using (C2), we get $\mathbf{a}_0 > 0$. Furthermore, using (6) and $\mu_{e,0} > \mu_{w,0}$ (by [C1]), we obtain $(1-\delta)\alpha_{e,0} - \mathbf{a}_0\gamma_{e,0} > 0$ such that $\mathbf{t}_g^+ > 0$. Hence, we get $\mathbf{t}_g \in (0, \mathbf{t}_g^+]$. Now consider the behavioral rule. The behavioral rule has an incentive to switch to wage work when $b_{e,t} \leq b_{w,0}$. As $\mu_{e,t} + \lambda\sigma_e/\gamma_{e,t} \leq \mu_{e,t} + \lambda\sigma_{e,t}/\gamma_{e,t}$, where $\gamma_{e,t} = \gamma_{e,0} + t$, a transition incentive will occur later than period $t = \mathbf{t}_b^-$, where \mathbf{t}_b^- solves $\mu_{e,t} + \lambda\sigma_e/\gamma_{e,t} = \mu_{w,0} + \lambda\sigma_{w,0}/\gamma_{w,0}$. Hence, we have $\mathbf{t}_b \geq \mathbf{t}_b^-$, where

$$\mathbf{t}_b^- = \frac{\mathbf{a}_1\lambda + \mathbf{a}_2}{\mathbf{a}_3\lambda + \mathbf{a}_4}$$

$$\mathbf{a}_1 := \gamma_b [\gamma_{w,0}\sigma_e - \gamma_{e,0}\sigma_{w,0}], \quad \mathbf{a}_2 := \alpha_{e,0}\gamma_{w,0} - \alpha_{w,0}\gamma_{e,0}$$

¹⁷Strictly speaking, this does not establish the existence of \mathbf{t}_b and \mathbf{t}_g . However, it is easy to demonstrate that when $t \rightarrow \infty$, such that the signal approaches $\mathbf{s}_{e,\infty}^*$, $g_{e,t}$, $b_{e,t} \rightarrow 0$. Given $g_{w,0}$, $b_{w,0} > 0$, a transition is guaranteed.

$$\mathbf{a}_3 := \gamma_b \sigma_{w,0} > 0, \quad \mathbf{a}_4 := \alpha_{w,0} > 0$$

It is obvious that $\mathbf{a}_3 \lambda + \mathbf{a}_4 > 0$, as, according to condition (C3), $\lambda > -\mathbf{a}_4/\mathbf{a}_3 = -\alpha_{w,0}[\gamma_b \sigma_{w,0}]^{-1}$. According to (C2) and (C4), we have $\mathbf{a}_1 < 0$; and, according to (C4), we have $\mathbf{a}_2 > 0$. Hence: $\mathbf{a}_1 \lambda + \mathbf{a}_2 > 0$. Thus, we have $\mathbf{t}_b^- > 0$ and $\mathbf{t}_b \in [\mathbf{t}_b^-, \infty)$. Now consider

$$\mathbf{t}_b^- - \mathbf{t}_g^+ = \frac{\mathbf{a}_5 \lambda + \mathbf{a}_6}{\mathbf{a}_0(\mathbf{a}_3 \lambda + \mathbf{a}_4)}$$

$$\mathbf{a}_5 := \mathbf{a}_0 \mathbf{a}_1 - (1 - \delta) \mathbf{a}_3 \alpha_{e,0} + \mathbf{a}_0 \mathbf{a}_3 \gamma_{e,0}$$

$$\mathbf{a}_6 := \mathbf{a}_0 \mathbf{a}_2 - (1 - \delta) \mathbf{a}_4 \alpha_{e,0} + \mathbf{a}_0 \mathbf{a}_4 \gamma_{e,0}$$

Given that $\lambda > -\mathbf{a}_4/\mathbf{a}_3$, we have $\mathbf{a}_0(\mathbf{a}_3 \lambda + \mathbf{a}_4) > 0$. Note that

$$\mathbf{a}_5 = \gamma_b [\delta(\alpha_{e,0} \sigma_{w,0} - \alpha_{w,0} \sigma_e) - (\alpha_{e,0} \sigma_{w,0} - \alpha_{w,0} \sigma_e) - \delta \gamma_{w,0} \sigma_e^2]$$

$$\mathbf{a}_6 = -\delta \alpha_{e,0} \gamma_{w,0} \sigma_e < 0$$

where $\delta < 1$. According to (C2) and (C4): $\alpha_{e,0} \sigma_{w,0} > \alpha_{w,0} \sigma_e$ such that $\mathbf{a}_5 < 0$. Consequently, if

$$\lambda < -\frac{\mathbf{a}_6}{\mathbf{a}_5} < 0 \quad (18)$$

we have $\mathbf{t}_b^- - \mathbf{t}_g^+ > 0$. To see that (18) is consistent with (C3), note that if $\mathbf{a}_7 := \alpha_{w,0} \mathbf{a}_5 - \gamma_b \sigma_{w,0} \mathbf{a}_6 < 0$, we have $-\mathbf{a}_6/\mathbf{a}_5 > -\alpha_{w,0}[\gamma_b \sigma_{w,0}]^{-1}$ and (18) does not require uncertainty aversion exceeding the boundary in (C3). Notice that

$$\begin{aligned} \mathbf{a}_7 &= \delta \gamma_b (\alpha_{e,0} [\alpha_{w,0} \sigma_{w,0} + \gamma_{w,0} \sigma_e] - \alpha_{w,0}^2 \sigma_e - \alpha_{w,0} \sigma_{w,0} \sigma_e^2) \\ &\quad + \alpha_{w,0} \gamma_b (\alpha_{w,0} \sigma_e - \alpha_{e,0} \sigma_{w,0}) \end{aligned}$$

where $\mathbf{a}_7 < 0$ requires

$$\alpha_{e,0} > \frac{\sigma_{w,0} \sigma_e - (1 - \delta) \gamma_b \alpha_{w,0}}{\gamma_b [\delta \gamma_{w,0} \sigma_e - (1 - \delta) \alpha_{w,0} \sigma_{w,0}]} \alpha_{w,0} \sigma_e \quad (19)$$

i.e., the prior number of successes in entrepreneurship must be sufficiently large. Note that $\delta \gamma_{w,0} \sigma_e - (1 - \delta) \alpha_{w,0} \sigma_{w,0} < 0$ given that $\alpha_{w,0} > \delta \sigma_e [\delta \sigma_{w,0} - \delta \sigma_e + \sigma_{w,0}]^{-1} \beta_{w,0}$, where the latter is consistent with $\alpha_{w,0} < \beta_{w,0}$ (important for overconfidence with respect to entrepreneurship), as $\delta \sigma_e [\delta \sigma_{w,0} - \delta \sigma_e + \sigma_{w,0}]^{-1}$ is smaller than 1 if $2\delta \sigma_e < (1 - \delta) \sigma_{w,0}$ and decreases if σ_e decreases. To demonstrate the second part of the proposition, note that either $\mathbf{t}_b < t_\sigma$, or $\mathbf{t}_b = t_\sigma$, or $\mathbf{t}_b > t_\sigma$. $\mu_{e,t}$ always decreases in t (given signal $\mathbf{s}_{e,t}^*$) such that $\mu_{e,t+1} < \mu_{e,t}$. $\gamma_{e,t}$ always increases in t such that $\gamma_{e,t+1} > \gamma_{e,t}$. If $\mathbf{t}_b < t_\sigma$, uncertainty increases in t such that $\sigma_{e,t+\Delta_t} > \sigma_{e,t}$, where $\Delta_t > 0$ is sufficiently small. If $\mathbf{t}_b = t_\sigma$, $\sigma_{e,t_\sigma-\Delta_t} < \sigma_{e,t_\sigma}$ and $\sigma_{e,t_\sigma+\Delta_t} < \sigma_{e,t_\sigma}$. If $\mathbf{t}_b > t_\sigma$, uncertainty decreases in t such that $\sigma_{e,t+\Delta_t} < \sigma_{e,t}$. Let $\mathbf{t}_b(\lambda)$ such that $\mathbf{b}_{e,\mathbf{t}_b(\lambda)} = \mathbf{b}_{w,0}$ or

$$\Delta_{\mu,\mathbf{t}_b(\lambda)} = \lambda \Delta_{\sigma,\mathbf{t}_b(\lambda)} \quad (20)$$

$$\Delta_{\mu, \mathbf{t}_b(\lambda)} := \mu_{e, \mathbf{t}_b(\lambda)} - \mu_{w, 0}, \quad \Delta_{\sigma, \mathbf{t}_b(\lambda)} := \gamma_b \left[\frac{\sigma_{w, 0}}{\gamma_{w, 0}} - \frac{\sigma_{e, \mathbf{t}_b(\lambda)}}{\gamma_{e, \mathbf{t}_b(\lambda)}} \right]$$

The right hand side of (20) is negative by (C2). Hence, we must have $\mu_{e, \mathbf{t}_b(\lambda)} < \mu_{w, 0}$, respectively $\Delta_{\mu, \mathbf{t}_b(\lambda)} < 0$. Let $\lambda' < \lambda$. Consequently:

$$\Delta_{\mu, \mathbf{t}_b(\lambda)} > \lambda' \Delta_{\sigma, \mathbf{t}_b(\lambda)} \quad (21)$$

Now consider how $\mathbf{t}_b(\lambda)$ must change if we want to establish an equality instead of the inequality in (21), i.e., if we want a $\mathbf{t}_b(\lambda')$ solving $\Delta_{\mu, \mathbf{t}_b(\lambda')} = \lambda' \Delta_{\sigma, \mathbf{t}_b(\lambda')}$. Assume that $\mathbf{t}_b(\lambda) < t_\sigma$. Here, it is obvious that it does not make sense to decrease \mathbf{t}_b , as, in such a case, $\Delta_{\mu, \mathbf{t}_b}$ would increase, while $\lambda' \Delta_{\sigma, \mathbf{t}_b}$ would decrease. Hence, we must have $\mathbf{t}_b(\lambda') > \mathbf{t}_b(\lambda)$ for $\mathbf{t}_b(\lambda) < t_\sigma$. Assume that $\mathbf{t}_b(\lambda) > t_\sigma$. In such a case, it is required that $\mu_{e, t}$ decreases faster over time than $\lambda \gamma_b \sigma_{e, t} / \gamma_{e, t}$ if $t = \mathbf{t}_b(\lambda)$ is the *first* period where $b_{e, t} = b_{w, 0}$. In the contrary case, $t = \mathbf{t}_b(\lambda)$ cannot be the first period solving $b_{e, t} = b_{w, 0}$. Hence, here again, we must increase \mathbf{t}_b such that $\Delta_{\mu, \mathbf{t}_b}$ decreases, while $\lambda' \Delta_{\sigma, \mathbf{t}_b}$ also decreases, however, not as fast as $\Delta_{\mu, \mathbf{t}_b}$. (If we decrease \mathbf{t}_b , $\Delta_{\mu, \mathbf{t}_b}$ increases, whereas $\lambda' \Delta_{\sigma, \mathbf{t}_b}$ also increases but slower than $\Delta_{\mu, \mathbf{t}_b}$.) Consequently, we must have $\mathbf{t}_b(\lambda') > \mathbf{t}_b(\lambda)$ for $\mathbf{t}_b(\lambda) > t_\sigma$. If we assume that $\mathbf{t}_b(\lambda) = t_\sigma$ and decrease \mathbf{t}_b , we have the case $\mathbf{t}_b(\lambda) < t_\sigma$, where decreasing \mathbf{t}_b cannot balance (21) but an increase in \mathbf{t}_b can; and if we assume that $\mathbf{t}_b(\lambda) = t_\sigma$ and increase \mathbf{t}_b , we have the case $\mathbf{t}_b(\lambda) > t_\sigma$, where increasing \mathbf{t}_b can balance (21) but not decreasing. To summarize, in all considered cases, we must have $\mathbf{t}_b(\lambda') > \mathbf{t}_b(\lambda)$. ■

Proposition 5 provides two important insights.¹⁸ First, individuals with sufficiently high uncertainty aversion will correct their false beliefs less efficiently than optimally-behaving individuals. Second, higher uncertainty aversion reduces the efficiency of the correction of false beliefs. The general logic behind the latter result is as follows. Assume that we have two individuals, *A* and *B*, with exactly the same initial prior observations, $\mathbb{x}_{e, 0}$ and $\mathbb{x}_{w, 0}$, but one individual, say *B*, is more uncertainty-averse than the other. If both individuals observe signal $\mathbf{s}_{e, t}^*$, their expected success probabilities will decrease by exactly the same amount. So, type-1 overconfidence is reduced by the same degree. However, due to type-2 overconfidence, the more uncertainty-averse individual *B* will transition to wage work later than *A*. The reason behind is that, given type-2 overconfidence leading to lower uncertainty in entrepreneurship than in wage work, the more uncertainty-averse individual *B* dislikes the fact that wage-work success is relatively more uncertain than entrepreneurial success more than *A*. Consequently, a larger number of non-successes in entrepreneurship is required to convince the more uncertainty-averse

¹⁸The conditions in Proposition 5, (C1)–(C4), are not especially restrictive and they are intuitively reasonable. For example, (C1) is needed to ensure that the Gittins index selects entrepreneurship in the first period. (C3) secures that wage work, and not unemployment, is the best alternative to entrepreneurship. Most conditions can be ensured by enough prior success observations in entrepreneurship and a sufficiently large amount of information: $\gamma(\mathbb{x}_{e, 0})$ and $\gamma(\mathbb{x}_{w, 0})$ must be both sufficiently large. In the third version of the model, it is demonstrated that effects predicted by Proposition 5 are present even if some conditions in the proposition are violated.

individual B to transition to wage work than the less uncertainty-averse individual A .

Also, Proposition 5 delivers a simple explanation for the earnings puzzle, the permanent selection of entrepreneurship although wage work is the better option. As high uncertainty aversion is empirically likely (see Section 2) and individuals with a sufficiently high level of uncertainty aversion need a long period of time to correct overconfidence with respect to entrepreneurship, observing some finite time interval of an individual's career can lead to the conclusion that she permanently selects an non-optimal option. Furthermore, under realistic conditions, the signal will be more ambiguous than the clear signal in Proposition 5—in reality, some periods in entrepreneurship might be successful. This will further reduce the efficiency of belief correction (demonstrated in Subsection 5.3).

5.2 Choices in model with predetermined absorption

Let us examine the second version of the bandit choice model: a model with predetermined absorption, henceforth, denoted by PAM, which is especially suitable to discuss transitions and sampling from different options.¹⁹ PAM has the following setup:

PREDETERMINED ABSORPTION MODEL (PAM). Assume that an individual learns if she is successful in an occupation one period after selecting the occupation for the first time. Consequently, there is uncertainty about success probabilities, $\sigma_i > 0$ for all $i \in \mathbb{O}_{-u}$, until an option is selected once. However, an option selected for the second time is not uncertain anymore. The individual is successful in occupation $i \in \mathbb{O}_{-u}$, and receives $\pi_i = 1$ in all periods the occupation is selected, with probability $\phi_i \in (0, 1)$, and the individual is not successful, and receives $\pi_i = 0$, with probability $1 - \phi_i$. The expected success probability is $\mu_i \in (0, 1)$ for all $i \in \mathbb{O}_{-u}$, where we assume that $\mu_e \neq \mu_w$. The number of prior observations is γ_i for all $i \in \mathbb{O}_{-u}$. μ_i , γ_i , and σ_i are computed on the basis of prior observations \mathbb{x}_i , where it is assumed that $\sigma_w \neq \sigma_e$, $\gamma_w \neq \gamma_e$, and $\sigma_w/\gamma_w \neq \sigma_e/\gamma_e$. Lastly, assume that it is commonly known that an option selected twice *must be selected* in all consecutive periods, i.e., an option selected twice is absorbing.

Given this particular setting, the number of decision sequences is restricted, thus it is possible to illustrate when and why uncertainty-averse gamblers do not exhibit optimal behavior. First, notice the following:

LEMMA 1. *Let r_{i,h_i} denote the index of option i selected h_i times. Given PAM, we obtain the following Gittins indices:*

$$g_{u,h_u} = 0 \quad \text{for all } h_u \quad (22)$$

$$g_{i,0} = \frac{\mu_i}{\delta[\mu_i - 1] + 1} \quad \text{for all } i \in \mathbb{O}_{-u} \quad (23)$$

¹⁹The approach is inspired by Berninghaus & Seifert-Vogt (1987) who use a similar method in a migration model.

or equivalently in behavioral-index-type notation

$$g_{i,0} = \mu_i + \omega_{g,i}\sigma_i \quad \text{for all } i \in \mathbb{O}_{-u} \quad (24)$$

where

$$\omega_{g,i} := \frac{\mu_i[1 - \mu_i]\delta}{\sigma_i[\delta(\mu_i - 1) + 1]} > 0$$

such that $g_{i,0} > \mu_i$ for $\delta \in (0, 1)$. For the second time an option $i \in \mathbb{O}_{-u}$ is selected, we obtain

$$g_{i,h_i} = \pi_i \in \{0, 1\} \quad \text{for } h_i > 0 \quad (25)$$

Proof. Equation (22) follows because unemployment is not uncertain and the unemployment reward is always zero. Equation (25) is a consequence of the PAM setting as an option selected once is not uncertain anymore and the gambler receives $\pi_i = 1$ or $\pi_i = 0$ for all remaining periods. To derive (23), use the optimal-stopping formulation of the Gittins index (see Equation [9]). First, note that it does not make sense to stop before observing the reward π_i at least once. After observing the reward, we either continue to play option i if $\pi_i = 1$ and obtain $\sum_{t=1}^{\infty} \delta^{t-1} \pi_i = 1/(1 - \delta)$, which, from the point of view of the gambler, occurs with probability μ_i , or we stop after observing the reward once if $\pi_i = 0$ and obtain zero, which, from the perspective of the gambler, occurs with probability $1 - \mu_i$. By dividing by the expected sum of discounted time, we get

$$g_{i,0} = \frac{\mu_i(1 - \delta)^{-1}}{\mu_i(1 - \delta)^{-1} + 1 - \mu_i} \quad \text{for all } i \in \mathbb{O}_{-u}$$

which corresponds to (23). ■

In contrast to the full model, the Gittins index of an initially uncertain option can become zero in the PAM setting as an option selected twice is not uncertain anymore. Consequently, given PAM, a gambler following the Gittins rule might at some point become indifferent between unemployment and all available non-unemployment options. Let us assume that in such a case the gambler will select unemployment, as unemployment is a minimal-effort option. Furthermore, a Gittins-index-based strategy is optimal in PAM:

LEMMA 2. *A Gittins-index strategy is optimal in PAM conditional on prior information \times_w and \times_e .*

Proof. The optimal deterministic strategy in PAM is relatively intuitive and easy to prove. Let OSPAM denote the optimal strategy in PAM. OSPAM prescribes the following choices: Let $i^{(1)} := \arg \max_{i \in \mathbb{O}_{-u}} \mu_i$. In the first period, select the option with the largest expected success probability $\mu_{i^{(1)}} = \max\{\mu_e, \mu_w\}$. Learn everything about the reward of $i^{(1)}$. If $\pi_{i^{(1)}} = 1$, continue to select option $i^{(1)}$ forever and obtain $1/(1 - \delta)$, which is the best possible outcome. If $\pi_{i^{(1)}} = 0$, switch to the option with the second-largest expected success probability $i^{(2)} := \mathbb{O}_{-u} \setminus i^{(1)}$ and observe the reward of $i^{(2)}$. If $\pi_{i^{(2)}} = 1$, continue to play $i^{(2)}$ and obtain $\delta/(1 - \delta)$,

which is the second-best possible outcome. If $\pi_{i^{(2)}} = 0$, all options have the same reward of zero such that it does not matter which option is selected. OSPAM yields $\mathbb{E}[\sum_{t=1}^{\infty} \delta^{t-1} \Pi_t | \mathbb{X}_e, \mathbb{X}_w] = \mu_{i^{(1)}}/(1-\delta) + [1 - \mu_{i^{(1)}}] \mu_{i^{(2)}} \delta / (1-\delta)$, which cannot be improved by any other deterministic strategy.²⁰ The Gittins index prescribes a strategy equivalent to OSPAM: In the first period, we have $g_{i,0} > g_{u,h_u} = 0$ for all $i \in \mathbb{O}_{-u}$ and all h_u ; and $g_{i,0} > g_{j,0}$ if $\mu_i > \mu_j$ or $g_{j,0} > g_{i,0}$ if $\mu_j > \mu_i$ for $i, j \in \mathbb{O}_{-u}$ where $i \neq j$. Hence, $i^{(1)}$ is selected. In the second period, we have $g_{i,1} > g_{j,0}$ for $i, j \in \mathbb{O}_{-u}$ if $\pi_i = 1$ and $g_{i,1} < g_{j,0}$ if $\pi_i = 0$ for $i, j \in \mathbb{O}_{-u}$ where $i \neq j$. Hence, option $i^{(1)}$ is selected if $\pi_{i^{(1)}} = 1$ and $i^{(2)}$ is selected if $\pi_{i^{(1)}} = 0$, and so forth. ■

The behavioral rule sets the following indices:

$$b_{i,0} = \mu_i + \lambda \gamma_b \frac{\sigma_i}{\gamma_i} \quad \text{for all } i \in \mathbb{O}_{-u} \quad (26)$$

$$b_{i,h_i} = \pi_i \in \{0, 1\} \quad \text{for all } h_i > 0 \quad (27)$$

for a non-unemployment option $i \in \mathbb{O}_{-u}$ and

$$b_{u,h_u} = 0 \quad \text{for all } h_u \quad (28)$$

for unemployment.

Using Lemma 1 and behavioral-rule indices (26)–(28), we can derive all plausible decision sequences in PAM. Table 1 compiles all decision sequences generated by the optimal or the behavioral rule—or both. For instance, sequence (DS4) should be read as follows. In the first period, the gambler selects wage work. In the second period, the gambler transitions to entrepreneurship. In the third period, the gambler selects entrepreneurship again. As an option selected twice must be selected in all consecutive periods, entrepreneurship is selected in all periods $t > 3$. There are only 9 plausible sequences. Some of the sequences, like (DS1) and (DS6), are produced by the behavioral rule only. Table 1 is used as an aid for proofs.

The two most interesting decision aspects most easily examined in the PAM setting are transition intensity and sampling frequencies. Transition intensity is defined as follows.

DEFINITION 5. (Transition intensity) Let $E_{r,t} := \{d_{r,t} \neq d_{r,t-1} \text{ for } t > 1\}$ denote a transition event generated by rule r . Transition intensity is defined as the ratio of the number of actual transitions to the number of all plausible transitions:

$$\hat{i}_r := \frac{\sum_{t=2}^{\infty} \mathbb{1}\{E_{r,t}\}}{\text{Maximal number of plausible transitions}} \in [0, 1] \quad (29)$$

²⁰To see, note, first, that it does not make sense to start with option $i^{(2)}$ or u , as both yield less than $\mu_{i^{(1)}}/(1-\delta)$. If $\pi_{i^{(1)}} = 0$, it does not make sense to switch to option u or to continue with $i^{(1)}$, as both actions yield zero. Consequently, the gambler should switch to $i^{(2)}$, as $\mu_{i^{(2)}} \delta / (1-\delta) > 0$. If $\pi_{i^{(1)}} = 1$, it does not make sense to switch from $i^{(1)}$, as $i^{(1)}$ yields $\delta / (1-\delta)$, while $i^{(2)}$ and u both generate a smaller expected reward— $i^{(2)}$ generates $\mu_{i^{(2)}} \delta / (1-\delta) < \delta / (1-\delta)$ (as $\mu_{i^{(2)}} < 1$) and u generates zero.

TABLE 1: Decision sequences (DS) plausible in PAM

	DECISION SEQUENCE	NUMBER OF TRANSITIONS	GENERATED BY WHICH RULE
(DS1)	$d_1 = u \rightarrow d_2 = u \rightarrow \dots$	0	b but not g
(DS2)	$d_1 = w \rightarrow d_2 = w \rightarrow \dots$	0	b and g
(DS3)	$d_1 = e \rightarrow d_2 = e \rightarrow \dots$	0	b and g
(DS4)	$d_1 = w \rightarrow d_2 = e \rightarrow d_3 = e \rightarrow \dots$	1	b and g
(DS5)	$d_1 = e \rightarrow d_2 = w \rightarrow d_3 = w \rightarrow \dots$	1	b and g
(DS6)	$d_1 = w \rightarrow d_2 = u \rightarrow d_3 = u \rightarrow \dots$	1	b but not g
(DS7)	$d_1 = e \rightarrow d_2 = u \rightarrow d_3 = u \rightarrow \dots$	1	b but not g
(DS8)	$d_1 = w \rightarrow d_2 = e \rightarrow d_3 = u \rightarrow d_4 = u \rightarrow \dots$	2	b and g
(DS9)	$d_1 = e \rightarrow d_2 = w \rightarrow d_3 = u \rightarrow d_4 = u \rightarrow \dots$	2	b and g

NOTE: The interpretation of, for instance, (DS9) is that entrepreneurship is selected in the first period, wage work is selected in the second period, unemployment is selected in the third period, and unemployment is selected in the fourth period. If an option is selected for the second time (e.g., unemployment in the fourth period in [DS9]), it must be selected forever.

In PAM, contrary to the full model, the number of transitions is *a priori* finite by construction. Hence, we have $\sum_{t=2}^{\infty} \mathbb{1}\{E_{r,t}\} < \infty$ and \mathfrak{i}_r is well-defined. Option-specific sampling frequencies are defined as follows.

DEFINITION 6. (Sampling frequency) Sampling frequency is the number of times rule r samples from option i :

$$\mathfrak{f}_{r,i} := \sum_{t=1}^{\infty} \mathbb{1}\{d_{r,t} = i\} \quad (30)$$

It is obvious that there exists only one $i_r^\infty \in \mathbb{O}$ such that $\mathfrak{f}_{r,i_r^\infty} = \infty$. For all $i \neq i_r^\infty$, $\mathfrak{f}_{r,i}$ is finite.

The PAM setting allows for the derivation of two central results on the relation between uncertainty aversion and transition intensity, and on the relation between uncertainty aversion and sampling frequencies. The first result is that uncertainty aversion leads to suboptimal transition intensity:

PROPOSITION 6. (Uncertainty aversion and transition intensity) *Let \mathfrak{i}_g^{\max} denote the maximal level of transition intensity generated by rule g. Moreover, let $\mathfrak{i}_b^{\max}(\lambda)$ denote the maximal level of transition intensity generated by rule b given uncertainty-aversion level λ . There exists a sufficiently high level of uncertainty aversion such that the behavioral rule transitions less than optimal or*

$$\mathfrak{i}_g^{\max} > \mathfrak{i}_b^{\max}(\lambda)$$

Furthermore, an increase in uncertainty aversion can further restrict transitions such that

$$\mathbb{i}_b^{\max}(\lambda) > \mathbb{i}_b^{\max}(\lambda - \lambda_\Delta)$$

where λ is such that $\mathbb{i}_b^{\max}(\lambda) > 0$, $\lambda_\Delta > 0$, and λ_Δ is sufficiently large.

Proof. The proposition can be proven with the aid of Table 1. Notice that the maximal number of plausible transitions is 2 (see Table 1) such that $\mathbb{i}_t = 2^{-1} \sum_{t=2}^{\infty} \mathbb{1}\{E_{t,t}\}$. The Gittins index generates transition intensity $\mathbb{i}_g \in \mathbb{I}_g = \{0, 1/2, 1\}$. Hence, we have $\mathbb{i}_g^{\max} = \max \mathbb{I}_g = 1$. Let $\lambda > \lambda_1 := \max\{-\mu_e \gamma_e [\gamma_b \sigma_e]^{-1}, -\mu_w \gamma_w [\gamma_b \sigma_w]^{-1}\}$, where the maximum is unique since $\mu_e \neq \mu_w$ and $\sigma_e/\gamma_e \neq \sigma_w/\gamma_w$. Note that $\lambda > \lambda_1$ is consistent with uncertainty aversion as $\lambda_1 < 0$. If the behavioral rule starts with option e , it will generate either sequence (DS3), or (DS5), or (DS9). If the behavioral rule starts with option w , it will generate either sequence (DS2), or (DS4), or (DS8). The behavioral rule will never start with option u or select u in the second period if $\lambda > \lambda_1$. Consequently, we have $\mathbb{i}_b \in \mathbb{I}_{b, \lambda > \lambda_1} = \{0, 1/2, 1\}$ such that $\mathbb{i}_{b, \lambda > \lambda_1}^{\max} = \max \mathbb{I}_{b, \lambda > \lambda_1} = 1$. Let $\lambda_2 := \min\{-\mu_e \gamma_e [\gamma_b \sigma_e]^{-1}, -\mu_w \gamma_w [\gamma_b \sigma_w]^{-1}\}$. Assume that $\lambda_1 \geq \lambda > \lambda_2$. If the behavioral rule starts with option e , implying $\mu_e \gamma_e \sigma_e^{-1} > \mu_w \gamma_w \sigma_w^{-1}$, it will generate either sequence (DS3) or (DS7). If the behavioral rule starts with option w , implying $\mu_w \gamma_w \sigma_w^{-1} > \mu_e \gamma_e \sigma_e^{-1}$, it will generate either sequence (DS2) or (DS6). The behavioral rule will not start with option u if $\lambda > \lambda_2$. As a consequence, we get $\mathbb{i}_b \in \mathbb{I}_{b, \lambda_1 \geq \lambda > \lambda_2} = \{0, 1/2\}$ such that $\mathbb{i}_{b, \lambda_1 \geq \lambda > \lambda_2}^{\max} = \max \mathbb{I}_{b, \lambda_1 \geq \lambda > \lambda_2} = 1/2$. Assume that $\lambda \leq \lambda_2$. The only decision sequence the behavioral rule can generate is (DS1). Consequently, we get $\mathbb{i}_b \in \mathbb{I}_{b, \lambda \leq \lambda_2} = \{0\}$ such that $\mathbb{i}_{b, \lambda \leq \lambda_2}^{\max} = \max \mathbb{I}_{b, \lambda \leq \lambda_2} = 0$. Now, let $\lambda' \in (\lambda_1, 0)$, $\lambda'' \in (\lambda_2, \lambda_1]$, and $\lambda''' \in (-\infty, \lambda_2]$. It is obvious that $\lambda' > \lambda'' > \lambda'''$. Only λ' generates transition intensities such that $\mathbb{i}_b^{\max}(\lambda') = \mathbb{i}_g^{\max}$. For the remaining levels of uncertainty aversion, we get $\mathbb{i}_b^{\max}(\lambda'') < \mathbb{i}_g^{\max}$ and $\mathbb{i}_b^{\max}(\lambda''') < \mathbb{i}_g^{\max}$. Furthermore: $\mathbb{i}_b^{\max}(\lambda') > \mathbb{i}_b^{\max}(\lambda'') > \mathbb{i}_b^{\max}(\lambda''')$. ■

The link between uncertainty aversion and transition intensity established in Proposition 6—higher uncertainty aversion results in less transitions—is plausible as empirically observable uncertainty aversion is rather high and empirically observable transition intensity is low (see Section 2).

Secondly, it can be demonstrated that uncertainty aversion leads to oversampling from less uncertain options:

PROPOSITION 7. (Uncertainty aversion and sampling frequency) *Conditional on prior information \times_e and \times_w , and without loss of generality, assume that entrepreneurship is the most uncertain option with the least amount of information, i.e., $\sigma_e > \sigma_w$ and $\gamma_e < \gamma_w$. Let $\mathbb{f}_{b,i}(\lambda)$ denote the number of times the behavioral rule samples from option i given uncertainty-aversion level λ . Given a sufficiently high level of uncertainty aversion, the behavioral rule oversamples with a positive probability from options with less uncertainty than entrepreneurship such that*

$$\mathbb{P}\{\mathbb{f}_{b,w}(\lambda) > \mathbb{f}_{g,w} \vee \mathbb{f}_{b,u}(\lambda) > \mathbb{f}_{g,u}\} > 0$$

Proof. Assume that $\mu_e > \mu_w$. As a success in entrepreneurship is generated with probability ϕ_e and a success in wage work with probability ϕ_w , we obtain²¹

$$\mathbb{f}_{g,w} = \begin{cases} 0 & \text{with probability } \phi_e \\ 1 & \text{with probability } (1 - \phi_e)(1 - \phi_w) \\ \infty & \text{with probability } (1 - \phi_e)\phi_w \end{cases} \quad (31)$$

If we assume $\mu_w > \mu_e$, we obtain

$$\mathbb{f}_{g,w} = \begin{cases} 1 & \text{with probability } 1 - \phi_w \\ \infty & \text{with probability } \phi_w \end{cases} \quad (32)$$

For the sampling frequency of unemployment, we have

$$\mathbb{f}_{g,u} = \begin{cases} 0 & \text{with probability } 1 - (1 - \phi_e)(1 - \phi_w) \\ \infty & \text{with probability } (1 - \phi_e)(1 - \phi_w) \end{cases} \quad (33)$$

Equation (31), (32), and (33) combined yield $p_f := \mathbb{P}\{\mathbb{f}_{g,w} < \infty \vee \mathbb{f}_{g,u} < \infty\} > 0$. Let $\lambda < -\mu_e \gamma_e [\gamma_b \sigma_e]^{-1}$. Given this particular level of uncertainty aversion, we get $b_{e,0} < 0$. (*Case 1*) If $\mu_w > \mu_e (\sigma_w / \sigma_e) (\gamma_e / \gamma_w)$, where $(\sigma_w / \sigma_e) (\gamma_e / \gamma_w) < 1$, and $\lambda > -\mu_w \gamma_w [\gamma_b \sigma_w]^{-1}$, we get $b_{w,0} > 0$, $b_{w,h_w} > 0$ with probability ϕ_w , and $b_{w,h_w} = 0$ with probability $1 - \phi_w$, where $h_w > 0$. (*Case 2*) If $\mu_w > \mu_e (\sigma_w / \sigma_e) (\gamma_e / \gamma_w)$ but $\lambda < -\mu_w \gamma_w [\gamma_b \sigma_w]^{-1}$, we get $b_{w,0} < 0$. (*Case 3*) If $\mu_w < \mu_e (\sigma_w / \sigma_e) (\gamma_e / \gamma_w)$, we have $b_{w,0} < 0$. In case 1, we get either $\mathbb{f}_{b,w}(\lambda) = \infty$ and $\mathbb{f}_{b,u}(\lambda) = 0$ with probability ϕ_w or $\mathbb{f}_{b,w}(\lambda) = 1$ and $\mathbb{f}_{b,u}(\lambda) = \infty$ with probability $1 - \phi_w$. In case 2 and 3, we get $\mathbb{f}_{b,w}(\lambda) = 0$ and $\mathbb{f}_{b,u}(\lambda) = \infty$. Hence, there always exists a high-enough level of uncertainty aversion (a sufficiently small λ) such that $\mathbb{f}_{b,w}(\lambda) = \infty \vee \mathbb{f}_{b,u}(\lambda) = \infty$, which in combination with $p_f > 0$ yields the desired result. ■

5.3 Choices in model with unexpected termination

In a setting without predetermined absorption, it is not possible to use straightforward analytical examples to analyze decision sequences. However, it is possible to

²¹To derive (31), note that the Gittins index starts with option e if $\mu_e > \mu_w$. After selecting option e , the gambler learns everything about the option. With probability ϕ_e , the reward of entrepreneurship is 1 and the gambler generates decision sequence (DS3) in Table 1. In (DS3) only entrepreneurship is selected such that $\mathbb{f}_{g,w} = 0$. With probability $1 - \phi_e$, the reward of entrepreneurship is zero and the gambler selects wage work. With probability $(1 - \phi_e)(1 - \phi_w)$, the reward of entrepreneurship and wage work is zero and the gambler generates sequence (DS9) in Table 1. In (DS9) wage work is selected once such that $\mathbb{f}_{g,w} = 1$. With probability $(1 - \phi_e)\phi_w$, the reward of entrepreneurship is zero but the reward of wage work is 1. In such a case, the gambler generates decision sequence (DS5) in Table 1. In (DS5) wage work is selected for an infinite number of periods such that $\mathbb{f}_{g,w} = \infty$. All other sampling frequencies are obtained in a similar way.

numerically simulate the model given the assumption that termination is not expected. A model with unexpected termination, the third version of the bandit choice model, henceforth, denoted by UTM, stops after a fixed number of $T \in \mathbb{N}$ periods, while gamblers assume that they are playing a model with an infinite time horizon, such that they do not anticipate termination time T in any way. The resulting model can be specified as UTM- T . In effect, the model with unexpected termination time corresponds to observing a finite time interval of an infinitely played bandit game.

By simulating a UTM- T version of the bandit model, it is possible to build some intuition into model mechanics and to discuss uncertainty-aversion effects in a more descriptive way than before. Furthermore, a UTM- T allows for testing whether results derived in the model with predetermined absorption also hold in the full model. First, I introduce some assumptions necessary for simulations and two simulation scenarios. Then, I present and analyze choice patterns generated by the model.

5.3.1 Assumptions for simulations

General assumptions are as follows. Gamblers are assumed to be uncertainty averse. Unemployment never generates a success and this fact is common knowledge. Success probabilities of entrepreneurship and wage work are unknown. To simulate optimal decision behavior, I rely on an approximation of the Gittins index, denoted by $\hat{g}(\mathbf{x})$, provided by Brezzi & Lai (2002). The simulation proceeds as follows.²² In the first step, the algorithm generates rewards given some set of parameters ϕ_w and/or ϕ_e . In the second step, the algorithm applies the Gittins-index and the behavioral rule to data simulated in the first step and produces decision sequences given some set of parameters like uncertainty aversion, λ , priors, $\mathbf{x}_{e,0}$ and $\mathbf{x}_{w,0}$, etc. The two steps are sequentially repeated 10,000 times.

I employ two distinct scenarios to examine different aspects of occupational choices in an uncertain decision environment. In the first scenario, I examine the impact of uncertainty aversion on transition intensity (see Proposition 6) and sampling frequencies (see Proposition 7). The second scenario examines the effect of uncertainty aversion on the correction of overconfidence with respect to entrepreneurship (see Proposition 5). The first scenario is interesting because it illustrates the effects of uncertainty aversion on career choice probabilities and career dynamics (transitions). The second scenario provides an explanation for the earnings puzzle.

FIRST SCENARIO. Table 2 summarizes all assumption on model parameters that are necessary to examine the effect of uncertainty aversion on transition intensity and sampling frequencies. In the first scenario, I use a model version with 100 periods, i.e., UTM-100, where after $T = 100$ periods gamblers are forced to retire.

Proposition 7, making prediction with respect to uncertainty-aversion effects on sampling frequencies, assumes that one option is more uncertain than all alterna-

²²The simulation code is available upon request from the author.

TABLE 2: First scenario

PARAMETER	RELATION	ASSUMPTION
T	=	100
δ	=	0.97
λ	\in	$\{-20, -18, \dots, -4, -2\}$
γ^*	=	2
ϕ_e, ϕ_w	\in	$\{0.1, 0.3, 0.5, 0.7, 0.9\}$
$\alpha_{e,0}, \beta_{e,0}, \alpha_{w,0}, \beta_{w,0}$	\in	$\{5, 6\}$

tive options. Therefore, to analyze the effects of uncertainty aversion on sampling frequencies, I consider only those parameter combinations where one option is clearly more uncertain than all other options. By symmetry of assumptions in Table 2, we can assign the highest level of uncertainty aversion to either entrepreneurship or wage work without changing results. When examining sampling frequencies, I assume that entrepreneurship is the most uncertain option.

SECOND SCENARIO. The second scenario illustrates the problem of the correction of overconfidence with respect to entrepreneurship. Numerical assumptions are presented in Table 3. Given assumptions in Table 3, entrepreneurship is selected by

TABLE 3: Second scenario

PARAMETER	RELATION	ASSUMPTION
δ	=	0.97
λ	\in	$\{-3.0, -2.6, -2.2, -1.8, -1.4, -1.0, -0.6, -0.2\}$
γ^*	=	2
ϕ_e	\in	$\{0.05, 0.10, 0.15\}$
$(\alpha_{e,0}, \beta_{e,0})$	=	(10, 5)
$(\alpha_{w,0}, \beta_{w,0})$	=	(3, 6)

the Gittins index *and* the behavioral rule in the first period.²³ As the only variable of interest is the earliest period where individuals abandon entrepreneurship, \mathbf{t}_g for the optimal rule and \mathbf{t}_b for the behavioral rule, there is no fixed termination period. Parameter combinations are consistent with overconfidence with respect to entrepreneurship, as stated in Definition 4.²⁴

It is in fact rather difficult to provide a simulation scenario where the behavioral rule abandons entrepreneurship. As the behavioral index converges to ϕ_e ²⁵, we have to impose $\phi_e < b(\mathbf{x}_{w,0})$. For all $\phi_e > b(\mathbf{x}_{w,0})$ — $b(\mathbf{x}_{w,0})$ can become quite small as it decreases if uncertainty aversion increases—, the probability of a behavioral-index-induced correction of overconfidence is small. Furthermore, note that, besides

²³Assumptions result in $\hat{g}(\mathbf{x}_{e,0}) > \hat{g}(\mathbf{x}_{w,0})$, $b(\mathbf{x}_{e,0}) > b(\mathbf{x}_{w,0})$, and $b(\mathbf{x}_{w,0}) > 0$.

²⁴The expected success probability of entrepreneurship is larger than of wage work, $\mu(\mathbf{x}_{e,0} = (10, 5)) > \mu(\mathbf{x}_{w,0} = (3, 4))$; there is less uncertainty about success in entrepreneurship than in wage work, $\sigma(\mathbf{x}_{e,0} = (10, 5)) < \sigma(\mathbf{x}_{w,0} = (3, 4))$; and the actual probability to succeed in entrepreneurship is smaller than the expected one, $\phi_e < \mu(\mathbf{x}_{e,0} = (10, 5))$, where $\phi_e \in \{0.05, 0.10, 0.15\}$.

²⁵See Appendix A where (in the proof of [PB4]) it is demonstrated that $\mathbb{P}\{\lim_{h_i \rightarrow \infty} b(\mathbf{x}_{i,h_i}) = \phi_i\} = 1$, where h_i is the number of times option $i \in \mathbb{O}_{-u}$ is selected.

(C3) and (C4), conditions in Proposition 5 are violated. Consequently, the scenario exemplifies that a failure to efficiently correct overconfidence due to uncertainty aversion may occur even under conditions weaker than in Proposition 5. Lastly, notice that the signal used in the second scenario, $\mathbf{s}_{e,t}$ where each signal component is $\pi_{e,t} = 1$ with probability $\phi_e > 0$ or $\pi_{e,t} = 0$ with probability $1 - \phi_e > 0$, is more ambiguous, and, therefore, more realistic than the “complete failure signal” $\mathbf{s}_{e,t}^*$ used in Proposition 5.

5.3.2 Simulated decisions

Figure 3a depicts simulation averages of transition intensity, normalized by average optimal levels, conditional on different levels of uncertainty aversion. The figure reveals that uncertainty-averse individuals transition less than optimal and that the gap between optimal and uncertainty-aversion-affected transition intensity increases if uncertainty aversion increases. Very high uncertainty aversion results in no transitions at all. As the same result is obtained in a model with predetermined absorption, I conclude that uncertainty-aversion effects are consistent across all model versions discussed.

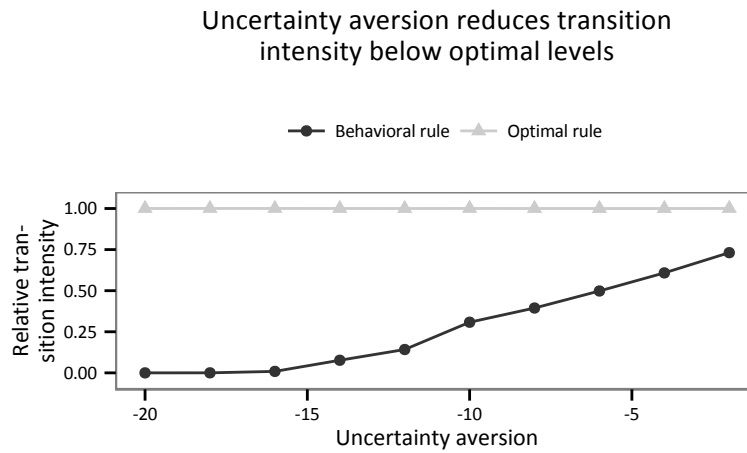
Figure 3a is consistent with observations on actual transitions (see Figure 1) and actual levels of uncertainty aversion (see Figure 2). In reality, high uncertainty aversion goes along with low transition intensity. In the model, increasing uncertainty aversion reduces transition intensity such that the aforementioned empirical result is reproduced.

In Figure 3b, it is assumed that entrepreneurship is the most uncertain option. To examine sampling from less uncertain options, the figure depicts simulation averages of sampling frequencies, normalized by the overall number of observations, of unemployment and wage work. The figure shows that high uncertainty aversion leads to oversampling from unemployment and wage work. Lower uncertainty aversion does not result in oversampling from unemployment but oversampling from wage work is still present. However, oversampling disappears if the level of uncertainty aversion is very small, such that individuals are almost uncertainty-neutral. The results are consistent with effects in a model with predetermined absorption. Insights may be summarized as follows.

SIMULATION RESULT 1. (Consistency between full model and model with predetermined absorption) *Empirically plausible uncertainty-aversion effects on transition intensity and effects on sampling frequencies derived inside the framework of a choice model with predetermined absorption are also present in a model with unexpected termination and, therefore, will also be present in the full model (observed over a finite time interval).*

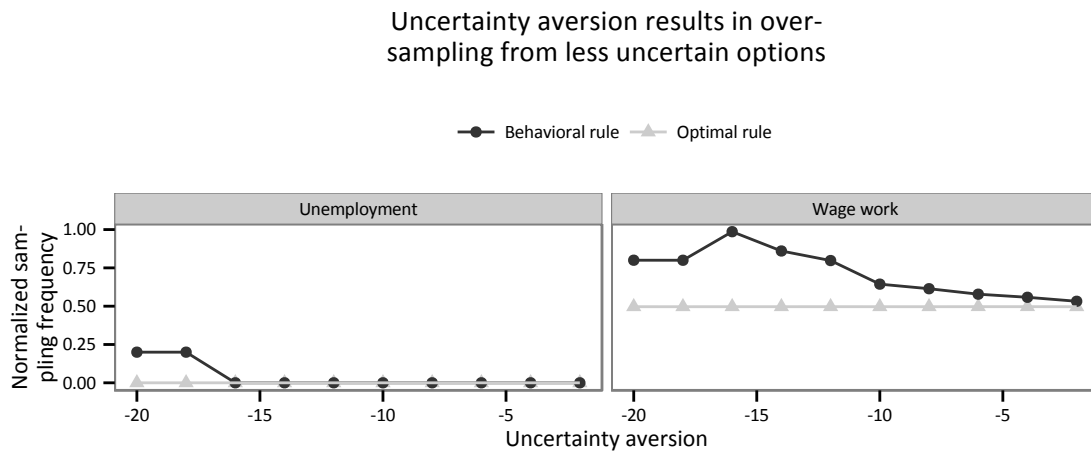
In Figure 4, I plot simulation averages of the earliest transition period (transitions from entrepreneurship to wage work), given that gamblers are overconfident with respect to entrepreneurship in early periods of their careers. In the figure, it is demonstrated that uncertainty-averse individuals transition later than optimal to

FIGURE 3: Transition intensity and sampling from less uncertain options (first simulated scenario)



NOTE: The figure presents simulation *averages* of transition intensity normalized by the optimal average level of transition intensity.

(a) Transition intensity



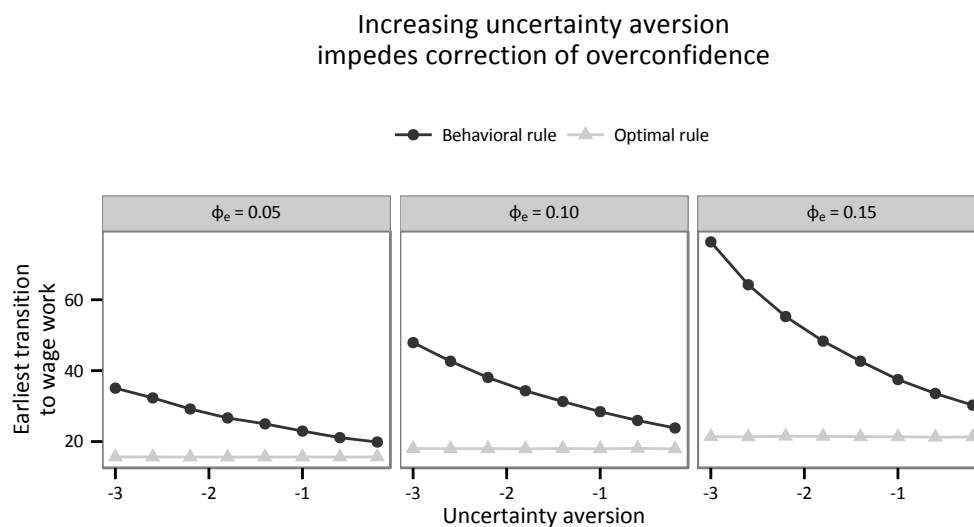
NOTE: The figure presents simulation *averages* of sampling frequencies. Frequencies are normalized by the overall number of observation, i.e., $T = 100$. It is assumed that entrepreneurship is the most uncertain option. Consequently, the figure presents sampling frequencies for wage work and unemployment, as these options are less uncertain than the most uncertain option.

(b) Oversampling from less uncertain options

wage work and that increasing uncertainty aversion delays the transition period, as predicted by Proposition 5.

An additional result, not part of Proposition 5, is that uncertainty-aversion-affected transition time is increasing in the “ambiguity” of the signal. To see, note that, given prior information $x_{e,0}$, the gambler rather expects a success than a non-success. If we consider a setting with $\phi_e = 0.05$, the signal the gambler receives is rather unambiguous as in most periods entrepreneurship generates a nonsuccess.

FIGURE 4: Correction of false beliefs (second simulated scenario)



NOTE: The figure presents simulation *averages* of the earliest transition period (from wage work to entrepreneurship).

In the setting with $\phi_e = 0.05$, an uncertainty-averse individual maximally needs less than 2.3 times longer than optimal to correct overconfidence. If we make the signal slightly more ambiguous by increasing the success probability of entrepreneurship from $\phi_e = 0.05$ to $\phi_e = 0.15$ (the signal predominantly generates nonsuccesses, however, there are more successes than in case of $\phi_e = 0.05$), an uncertainty-averse individual maximally needs more than 3.5 times longer than optimal to correct overconfidence with respect to entrepreneurship.

As there is plenty of evidence that entrepreneurs are overconfident (Busenitz & Barney 1997; Camerer & Lovo 1999; Bernardo & Welch 2001; Koellinger et al. 2007) and, according to the bandit choice model, uncertainty-averse and overconfident entrepreneurs learn inefficiently such that overconfidence may need an exceptionally long time to be corrected, the bandit choice model with a decision rule accounting for uncertainty aversion explains an observation like the earnings puzzle. Summarizing results of the second simulation scenario, we obtain the following:

SIMULATION RESULT 2. (Uncertainty aversion and overconfidence) *Higher uncertainty aversion is associated with a less efficient correction of false beliefs even when not all conditions in Proposition 5 hold. Furthermore, a signal more ambiguous than a complete failure may further impede the correction of overconfidence. In general, a non-optimal correction of overconfidence leading to earnings-puzzle-like observations is a likely scenario if entrepreneurs are uncertainty-averse.*

Summarizing the above, an increase in uncertainty aversion induces an increase in the propensity to select a relatively less uncertain option and impedes the correction of false beliefs. These predictions are tested with empirical data in the next section. Furthermore, the optimal rule and the behavioral rule (with a sufficiently low level of uncertainty aversion) produce transitions between occupations across all model versions, reproducing the empirical observation that many individuals have mixed careers consisting of entrepreneurship and wage work spells (see, e.g.,

Burke et al. 2008). In the bandit model, high uncertainty aversion is associated with a small number of transitions between occupations, which is consistent with empirical observations. Finally, as the capacity to correct false beliefs is limited by uncertainty aversion and entrepreneurs tend to be overconfident—they have too positive beliefs with respect to entrepreneurial success—, uncertainty aversion provides an explanation for the earnings puzzle.

6 Empirical evidence on model predictions

The bandit model generates a number of testable predictions. Two especially interesting predictions include, first, that higher uncertainty aversion will result in a higher probability to select a “sure” option; and, secondly, that higher uncertainty aversion will reduce the learning efficiency of entrepreneurs. Note that the two predictions hold at different levels: The first prediction applies to general occupational choices, while the second prediction is specific for entrepreneurs.

However, there is a significant problem with data availability. There are many data sets featuring a measure of risk aversion but there are few measures of uncertainty aversion. As the concepts of risk and uncertainty are quite different, risk aversion is not necessarily a good proxy for uncertainty aversion. The only established measure of uncertainty aversion that I am aware of is Hofstede’s uncertainty avoidance index (Hofstede et al. 2010), which was discussed in Section 2. Hofstede’s index is available for a number of countries but the index does not vary over time. Thus, relying on Hofstede’s index has the downside that one can only exploit the cross-country dimension where data availability is significantly restricted. Though the results could profit from more data (a replication study is definitely required), I demonstrate that both predictions above are consistent with available data.

6.1 First hypothesis: Uncertainty aversion increases propensity to avoid uncertain options

A testable version of the prediction that increasing uncertainty aversion will amplify the tendency to select a “sure” option, based on Proposition 4 and 7, is as follows.

HYPOTHESIS 1. (Avoiding uncertain options) Compared to a low level of uncertainty aversion, a higher aversion level is associated with a lower probability to select entrepreneurship because the propensity to select an option with low uncertainty increases.

In the bandit choice model, the low-uncertainty option is represented by unemployment. Some types of wage work can also be a low-uncertainty option. An example is Japan’s lifetime employment system where workers are hired immediately after their graduation and, on the basis of a gentlemen’s agreement, are not laid off until they retire (Sullivan & Peterson 1991). However, there is no indication that entrepreneurship generally represents the most certain option as entrepreneurs are usually confronted with new problems that cannot be fully anticipated, whereas

wage workers and the unemployed solve rather well-known problems. Furthermore, information about entrepreneurship is generally scarcer than information about non-entrepreneurial options (see, for instance, Kelley, Brush, Greene & Litovsky 2013).

6.1.1 Modeling uncertainty aversion and choice of entrepreneurship

I assess the first prediction with a clustering model. Assume that we can observe the uncertainty aversion and the probability to select entrepreneurship of individual (country) $j = 1, \dots, N_j$. Let $\mathbb{P}e_j$ denote j 's probability to select entrepreneurship and let, furthermore, λ_j^e denote j 's uncertainty-aversion type. Using a parsimonious model, I simply consider the joint distribution of $\mathbb{P}e_j$ and λ_j^e , where $\mathbb{P}e_j$ and λ_j^e are combined into vector $\mathbf{x}_j = [\lambda_j^e, \mathbb{P}e_j]^\top$. I assume that \mathbf{x}_j can be modeled with a Gaussian mixture model of the following form:

$$\prod_{j=1}^{N_j} \left\{ \sum_{c=1}^{N_c} p_c f_c(\mathbf{x}_j; \mathbf{m}_c, \Sigma_c) \right\} \quad (34)$$

$$f_c(\mathbf{x}_j; \mathbf{m}_c, \Sigma_c) = (2\pi)^{-1} |\Sigma_c|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} [\mathbf{x}_j - \mathbf{m}_c]^\top \Sigma_c^{-1} [\mathbf{x}_j - \mathbf{m}_c]\right)$$

N_j is the overall number of individuals or observations. N_c is the number of clusters. $p_c > 0$ is the probability that an observation belongs to cluster c . Note that $\sum_{c=1}^{N_c} p_c = 1$. $\mathbf{m}_c = [\bar{\lambda}_c, \overline{\mathbb{P}e}_c]^\top$ is a cluster-specific mean and Σ_c is a cluster-specific covariance matrix.

The mixture model in Equation (34) allows data to be clustered such that different clusters of observations might be generated by different distributions. The number of clusters is finite. Σ_c can be written in terms of its Eigen decomposition such that

$$\Sigma_c = \kappa_c \mathbf{A}_c \mathbf{S}_c \mathbf{A}_c^\top$$

where κ_c is the volume, \mathbf{A}_c the orientation, and \mathbf{S}_c the shape (Banfield & Raftery 1993). I use 10 different specifications of Σ_c that differ with respect to how they treat group differences. For instance, it can be assumed that volume and shape can be equal across clusters but not orientation. For the number of clusters, I assume a maximum of 9 clusters, i.e., $N_c \in \{1, 2, \dots, 9\}$.

In its most parsimonious version, Hypothesis 1 requires only two clusters: a cluster with moderate or low average uncertainty aversion and a cluster with high average uncertainty aversion. According to Hypothesis 1, the cluster with high average uncertainty aversion should exhibit a lower average probability to select entrepreneurship than the other cluster. I will show that given a number of alternative empirical models the best-performing model is the one suggested by Hypothesis 1.

6.1.2 Data on uncertainty aversion and choice of entrepreneurship

I use Hofstede's uncertainty avoidance index, available from Hofstede (2015) and described in Hofstede et al. (2010), as a proxy for uncertainty-aversion type. Hofstede's index is never negative and more uncertainty aversion is represented by higher index values. To make Hofstede's index and the bandit model's uncertainty preferences parameter λ comparable, I multiply the index by negative 1. After the transformation, more uncertainty aversion is represented by smaller index values.

To measure the probability to select entrepreneurship, I rely on established business ownership rates, representing the long-term tendency of a society to select entrepreneurship, provided by the Global Entrepreneurship Monitor and available from GEM (2015). Business ownership data is annual and the range of periods is 2001–2013 but not all periods are available for all countries. However, since Hofstede's index is time-invariant, I average over all available periods.²⁶ As a result, I have 67 joint observation of average business ownership rate and uncertainty-aversion type.

6.1.3 Joint distribution of uncertainty aversion and the probability to select entrepreneurship

Empirical models are estimated with the expectation-maximization algorithm.²⁷ Table 4 depicts the Bayesian information criterion (BIC) for estimated models. A

TABLE 4: Explanatory performance of different model specifications, where the best model is marked by an asterisk

NUMBER OF CLUSTERS	BIC OF BEST COVARIANCE SPECIFICATION
1	– 1006.12
2	– 997.92*
3	– 1003.24
4	– 1010.01
5	– 1019.69
> 5	≤ –1025.16

higher BIC indicates better explanatory performance given a penalty for model complexity.

According to Table 4, the best model has two clusters. The model suggested by

²⁶Data characteristics are provided in Appendix C.

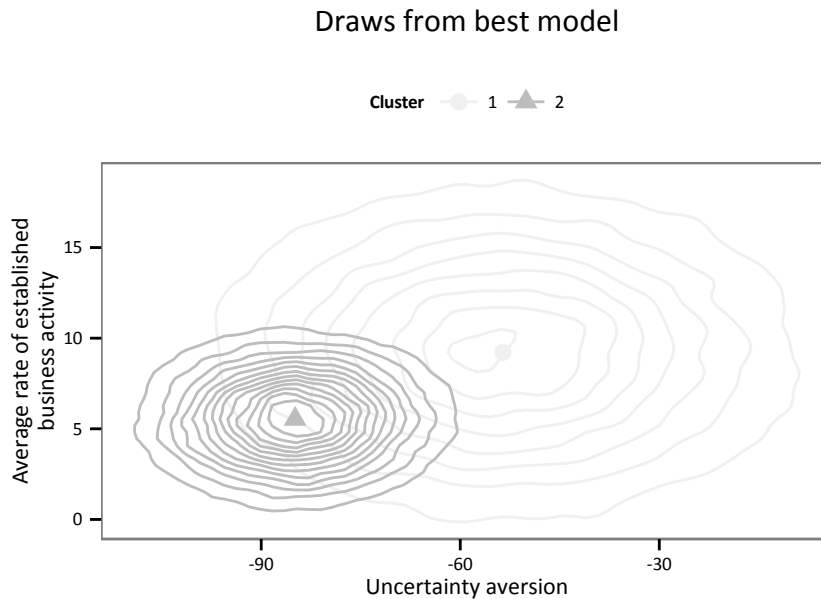
²⁷To estimate the model, I use the `mclust` package in R. An alternative suggested by Fraley & Raftery (2007) is to replace the maximum-likelihood estimator, in the expectation-maximization algorithm, by a maximum-posterior estimator. This modified approach can substantially reduce the number of non-fitted models. However, the general result (provided upon request) does not change. An additional concern is that established business activity may underestimate the amount of information about entrepreneurship (those who abandon entrepreneurship also provide information) such that the assumption that entrepreneurship is the most uncertain option available becomes problematic. To address this issue, I estimate the model only with countries that have an average business activity rate not above 10%. Results (provided upon request) are still consistent with Hypothesis 1.

Table 4 has the following cluster means:

$$\mathbf{m}_1 = \begin{bmatrix} \bar{\lambda}_1 = -54.297 \\ \bar{\mathbb{P}e}_1 = 9.129 \end{bmatrix}, \quad \mathbf{m}_2 = \begin{bmatrix} \bar{\lambda}_2 = -85.271 \\ \bar{\mathbb{P}e}_2 = 5.336 \end{bmatrix}$$

where the first element is average uncertainty-aversion type and the second the average rate of established business activity. Figure 5 shows 100,000 draws from

FIGURE 5: Density of the best model according to Bayesian information criterion



NOTE: Point and triangle indicate cluster means.

the best model's density. The average business ownership rate of the high-aversion cluster (the dark-gray cluster with mean \mathbf{m}_2) is lower than the corresponding rate of the cluster with low or moderate uncertainty aversion (the light-gray cluster with mean \mathbf{m}_1). Hence, I obtain the following result consistent with Hypothesis 1:

EMPIRICAL RESULT 1. (Avoiding uncertain options) *Higher uncertainty aversion is associated with a lower level of long-term entrepreneurial activity.*

6.2 Second hypothesis: Uncertainty aversion reduces learning efficiency of entrepreneurs

Consider two individuals whose only difference is how strong they dislike uncertainty. Both individuals exhibit the same level of overconfidence with respect to entrepreneurship. Assume that both individuals start an own company. According to the theoretical model, the more uncertainty-averse individual will learn less efficiently than the less uncertainty-averse individual. Consequently, the more uncertainty-averse individual has a *higher* probability to survive in entrepreneurship than the less uncertainty-averse individual given the same conditions because the less uncertainty-averse individual will correct her false beliefs (she is overconfident)

more efficiently than will the more uncertainty-averse individual (who is also overconfident). Hence, the bandit model makes the following testable prediction, on the basis of Proposition 5:

HYPOTHESIS 2. (Firm survival depends on uncertainty aversion) *Given that overconfidence of entrepreneurs is a common phenomenon, the probability of firm death should be lower for individuals with relatively high uncertainty aversion than for individuals with relatively low uncertainty aversion.*

6.2.1 Modeling uncertainty aversion and entrepreneurial survival

To test Hypothesis 2, I examine the impact of uncertainty aversion on firm death rates. Let $D \in (0, 1)$ denote the death rate of firms. Assume that death rates are individual- and time-specific, where individuals (countries in our setting) are indexed by j and time by t . Rates assume values in $(0, 1)$; the distribution of rates is not necessarily symmetric; and there also might be heteroskedasticity. Hence, a simple linear regression might not be appropriate. In such a setting, Ferrari & Cribari-Neto (2004) propose to use a beta regression.²⁸ Ferrari & Cribari-Neto (2004) suggest to parameterize the beta distribution in terms of its mean and precision.²⁹

Using the parameterization suggested by Ferrari & Cribari-Neto (2004), I assume the following:

$$D_{j,t} \sim \mathfrak{B}(\mu_{j,t}^D, \gamma^D), \quad \mu_{j,t}^D = \text{logit}^{-1}(\eta_{j,t}), \quad \gamma^D \text{ is constant} \quad (35)$$

such that death rates $D_{j,t}$ follow a beta distribution with mean $\mu_{j,t}^D$ and constant precision γ^D . $\eta_{j,t}$ is a linear function of uncertainty aversion type and a number of additional covariates that is connected to the mean by a logit-link function. Mean and variance both vary across time and individuals such that we do not have to assume homoskedasticity:

$$\mathbb{E}[D_{j,t}] = \mu_{j,t}^D, \quad \mathbb{V}[D_{j,t}] = \frac{\mu_{j,t}^D(1 - \mu_{j,t}^D)}{1 + \gamma^D}$$

Data is used to predict $\eta_{j,t}$ given by

$$\eta_{j,t} = \text{constant} + r_\lambda \mathbb{1}\{j \in H\} + \mathbf{c}_{j,t}^\top \mathbf{r} \quad (36)$$

The indicator $\mathbb{1}\{j \in H\}$, capturing uncertainty aversion types, is 1 if country j has relatively high uncertainty aversion (i.e., j belongs to the high-uncertainty-aversion group [type] H) and zero else. $\mathbf{c}_{j,t}$ is a vector of additional covariates. r_λ and \mathbf{r}

²⁸For further studies employing the relatively new beta regression approach, see, e.g., De Paola, Scoppa & Lombardo (2010) and Buntaine (2011).

²⁹Let μ^D denote the mean and γ^D the precision. A beta distribution is traditionally parameterized in terms of α^D and β^D such that density is $f(y; \alpha^D, \beta^D) = \text{B}(\alpha^D, \beta^D)^{-1} y^{\alpha^D-1} (1-y)^{\beta^D-1}$. A parameterization in terms of $\mu^D = \alpha^D / (\alpha^D + \beta^D)$ and $\gamma^D = \alpha^D + \beta^D$ yields density $f(y; \mu^D, \gamma^D) = \text{B}(\mu^D \gamma^D, [1 - \mu^D] \gamma^D)^{-1} y^{\mu^D \gamma^D-1} (1-y)^{[1-\mu^D] \gamma^D-1}$.

are coefficients to be estimated. According to Hypothesis 2, I expect $r_\lambda < 0$ such that belonging to the high-uncertainty-aversion type decreases (average) firm death rates.

As a robustness test, I also estimate a simple linear regression

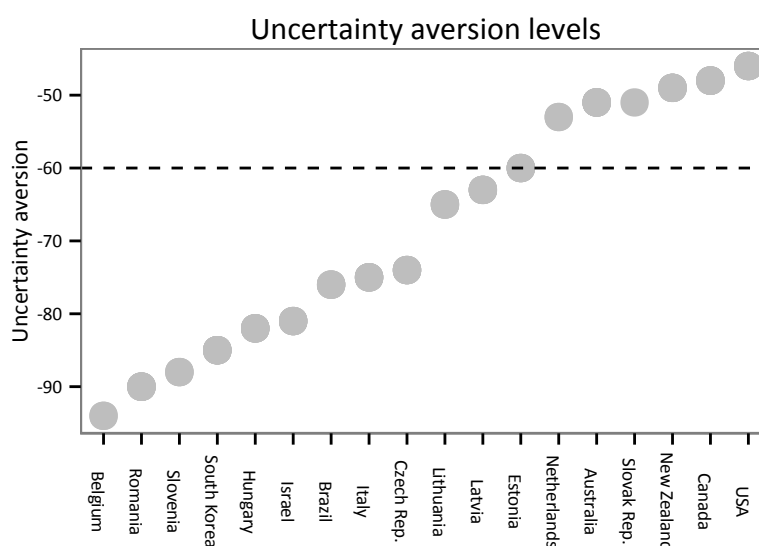
$$D_{j,t} = \eta_{j,t} + \varepsilon_D \quad (37)$$

where ε_D is an error term with the usual OLS properties.

6.2.2 Data on entrepreneurial survival

I use the same measure of uncertainty aversion as before, i.e., a transformed version of Hofstede's index. Figure 6 shows country-specific uncertainty aversion levels. I

FIGURE 6: Levels of uncertainty aversion (smaller values indicate more aversion)



assume that uncertainty aversion is relatively high if the level of uncertainty aversion is below -60 , which is indicated with a horizontal line in Figure 6.

OECD (2015) provides a small unbalanced panel of annual enterprise death rates for a number of countries. Death rates are reported for the period from 2006 to 2012 but some periods are missing for some countries.³⁰ Additional covariates (captured by vector \mathbf{c}) are the following: GDP growth, inflation, real interest rates, start-up costs, the unemployment rate, and average ease of doing business. All covariates are provided by the World Bank and available from World Bank (2015).

Average ease of doing business is constructed on the basis of annual indicators, where a higher indicator value corresponds to a situation where doing business is easier, and is used to measure between-country differences in ease of doing business.³¹ It should be noted that average ease of doing business and country-specific

³⁰Data characteristics of enterprise death rates are provided in Appendix C.

³¹In Appendix C, I provide characteristics of the ease of doing business indicator.

uncertainty aversion are highly correlated (correlation: 0.645; p -value: ≈ 0), implying that in countries where uncertainty aversion is lower doing business is easier (or *vice versa*). Hence, by including average ease of doing business indicators, it is possible to capture country differences that are related to the business environment but not directly related to uncertainty preferences.

The period 2006–2012 includes the financial crisis, which affected entrepreneurial survival. Therefore, I construct a crisis dummy that is 1 in year 2008, 2009, and 2010, but zero else.³² A total of 89 joint observations are available for analysis.

6.2.3 Effects of uncertainty aversion on the death of firms

The beta regression model, given by (35) and (36), is estimated by maximum likelihood with the bias correction proposed by Kosmidis & Firth (2010).³³ The linear model, given by (36) and (37), is estimated using OLS with heteroskedasticity-robust standard errors. Estimation results are presented in Table 5. With the exception of

TABLE 5: Determinants of firm death rates

VARIABLE	BETA REGRESSION		OLS	
	<i>Coefficient</i>	<i>Standard error</i>	<i>Coefficient</i>	<i>Standard error</i>
Constant	-4.094***	(0.639)	-0.047	(0.045)
Belonging to high- uncertainty-aversion group ^a	-0.256**	(0.120)	-0.016*	(0.009)
GDP growth	-0.025***	(0.009)	-0.002***	(0.001)
Inflation	0.036**	(0.014)	0.003***	(0.001)
Real interest rate	0.022***	(0.006)	0.002***	(0.000)
Start-up costs	0.041***	(0.008)	0.003***	(0.001)
Unemployment rate	0.007	(0.014)	0.001	(0.001)
Average ease of doing business ^b	0.019***	(0.007)	0.001***	(0.000)
Crisis dummy	0.139*	(0.083)	0.009	(0.006)
Precision	99.77***	(12.910)		
	Pseudo R ² = 0.352		R ² = 0.349	

^aIndicator is 1 if country belongs to high-uncertainty-aversion group and zero else.

^bAn increase in the indicator signifies that doing business becomes easier.

***significant at 1%-level; **significant at 5%-level; *significant at 10%-level

NOTE: Standard errors are in parentheses; in case of OLS, standard errors are robust.

significance levels and effect sizes, there is no substantial difference between the beta regression and OLS.

Significant effects not related to uncertainty aversion are all plausible—though, they are not necessarily causal. For instance, an increase in GDP growth significantly reduces death rates, which is a plausible effect since improved demand conditions

³²The crisis dummy is, as it might be expected, positively correlated with entrepreneurial death rates (correlation: 0.266) and the correlation is significant at a 5%-level (p -value: 0.012).

³³Maximum-likelihood based beta regressions are implemented in the R-package *betareg*.

should foster entrepreneurial survival. An increase in real interest rates, corresponding to tighter restrictions on borrowing money, significantly increases death rates. An increase in start-up costs, reducing available resources for entrepreneurial survival by imposing higher entry barriers, also significantly increases death rates. The financial crisis significantly increases death rates—this results holds for the beta regression only. An interesting finding is that if doing business becomes easier, death rates are increased, which is a plausible result if ease of doing business reflects competition effects: If doing business becomes easier, competition is intensified such that survival probabilities decrease.

What is most important, belonging to the high-uncertainty-aversion type significantly decreases death rates in the beta and the OLS model: The estimated coefficient of uncertainty aversion, in Table 5, is negative and significant at the 5%-level in case of the beta regression and at the 10%-level in case of the OLS model. Consequently, I establish the following result consistent with Hypothesis 2:

EMPIRICAL RESULT 2. (Firm survival depends on uncertainty aversion) *Higher uncertainty aversion is associated with a higher probability of firm survival.*

To summarize, two central choice patterns predicted by the bandit model are consistent with available data.

7 Conclusion

Uncertainty is an unavoidable feature of daily life. Career decisions are also confronted with uncertain rewards. Yet, practically all approaches modeling occupational choices in the literature discuss choices under the assumption that choice outcomes are either deterministic or risky. The aim of this paper is to take a renewed perspective on occupational choice by constructing a model where, in contrast to existing literature, it is assumed that some choice outcomes are uncertain—the distributions of rewards in entrepreneurship and wage work are unknown.

This paper provides a tractable and flexible formulation of a Bayesian multi-armed bandit model with three arms representing the occupational options wage work, entrepreneurship, and unemployment. Decisions are derived from either an optimal rule or a decision rule accounting for uncertainty aversion, where the rule accounting for uncertainty aversion is consistent with available experimental results on the behavior of actual individuals playing bandits. The model generates career trajectories ranging from pure wage workers and entrepreneurs to individuals mixing spells in wage work, entrepreneurship, and unemployment. As mixed careers, which are not generated by most contemporary occupational choice models, are common and likely to become more important as traditional career paths (a worker is employed in one firm for a long period of time) gradually become less frequent, the bandit model is a particularly appropriate approach to analyze career choices.

With respect to effects of uncertainty preferences in a decision environment with elements of uncertainty, I show that sufficiently high uncertainty aversion has a number of negative consequences such as insufficient transitions between occupa-

tional options and oversampling from less uncertain options, where the benchmark is individually optimal behavior. Furthermore, sufficiently high uncertainty aversion introduces inefficiencies into the correction of false beliefs of individuals exhibiting overconfidence. In conjunction with the observation that entrepreneurs tend to be overconfident, inefficient learning explains an observation like the earnings puzzle.

Regarding observable decision patterns, the bandit model predicts that high-aversion societies have lower entrepreneurship levels than low-aversion societies and that, under plausible conditions, firm death rates should be higher in low-aversion than in high-aversion societies, as the former learn more efficiently than the latter. Although there are some restrictions on data availability, existing data supports both predictions.

The model presented in this paper has some advantages over alternative occupational choice models. A main advantage of the model is that career trajectories are a natural feature of the dynamic solution of individual occupational choice problems and that learning efficiency is linked to a simple psychological variable in the form of uncertainty aversion. Still, there are also shortcomings requiring further research. For instance, in the model, I postulate occupation-specific probabilities to succeed but success probabilities should rather not be treated as primitives, making it necessary to introduce lower-level determinants of occupational successes.

References

- Anderson, C. (2001). *Behavioral Models of Strategies in Multi-armed Bandit Problems*. PhD thesis, California Institute of Technology, Pasadena.
- Antonovics, K. & Golan, L. (2012). Experimentation and Job Choice. *Journal of Labor Economics*, 30(2), 333–366.
- Arthur, M. (1994). The Boundaryless Career: A New Perspective for Organizational Inquiry. *Journal of Organizational Behavior*, 15(4), 295–306.
- Arum, R. & Müller, W. (2004). Self-employment Dynamics in Advanced Economies. In R. Arum & W. Müller (Eds.), *The Re-emergence of Self-employment* (pp. 1–35). Princeton: Princeton Univ. Press.
- Astebro, T. & Chen, J. (2014). The Entrepreneurial Earnings Puzzle. *Journal of Business Venturing*, 29(1), 88–105.
- Astebro, T., Chen, J., & Thompson, P. (2011). Stars and Misfits: Self-Employment and Labor Market Frictions. *Management Science*, 57(11), 1999–2017.
- Banfield, J. & Raftery, A. (1993). Model-Based Gaussian and Non-Gaussian Clustering. *Biometrics*, 49(3), 803–821.
- Bellman, R. (1956). A Problem in the Sequential Design of Experiments. *Sankhyā: The Indian Journal of Statistics*, 16(3/4), 221–229.

- Benz, M. & Frey, B. (2008a). Being Independent is a Great Thing. *Economica*, 75(298), 362—383.
- Benz, M. & Frey, B. (2008b). The Value of Doing What You Like. *Journal of Economic Behavior and Organization*, 68(3/4), 445—455.
- Bergemann, D. & Hege, U. (2005). The Financing of Innovation. *RAND Journal of Economics*, 36(4), 719—752.
- Bernardo, A. & Welch, I. (2001). On the Evolution of Overconfidence and Entrepreneurs. *Journal of Economics & Management Strategy*, 10(3), 301—330.
- Bernhardt, A., Morris, M., Handcock, M., & Scott, M. (2001). *Divergent Paths: Economic Mobility in the New American Labor Market*. New York: Russell Sage.
- Berninghaus, S. & Seifert-Vogt, H. (1987). International Migration under Incomplete Information. *Swiss Journal of Economics and Statistics*, 123(2), 199—218.
- Biemann, T., Zacher, H., & Feldman, D. (2012). Career Patterns: A Twenty-year Panel Study. *Journal of Vocational Behavior*, 81(2), 159—170.
- Blanchflower, D. (2000). Self-employment in OECD Countries. *Labour Economics*, 7(5), 471—505.
- Brezzi, M. & Lai, T. (2000). Incomplete Learning from Endogenous Data in Dynamic Allocation. *Econometrica*, 68(6), 1511—1516.
- Brezzi, M. & Lai, T. (2002). Optimal Learning and Experimentation in Bandit Problems. *Journal of Economic Dynamics and Control*, 27(1), 87—108.
- Buntaine, M. (2011). Does the Asian Development Bank Respond to Past Environmental Performance when Allocating Environmentally Risky Financing. *World Development*, 39(3), 336—350.
- Burke, A., FitzRoy, F., & Nolan, M. (2008). What Makes a Die-hard Entrepreneur? Beyond the 'Employee or Entrepreneur' Dichotomy. *Small Business Economics*, 31(2), 93—115.
- Busenitz, L. & Barney, J. (1997). Differences Between Entrepreneurs and Managers in Large Organizations. *Journal of Business Venturing*, 12(1), 9—30.
- Camerer, C. & Ho, T.-H. (1999). Experience-weighted Attraction Learning in Normal Form Games. *Econometrica*, 67(4), 827—874.
- Camerer, C. & Lovallo, D. (1999). Overconfidence and Excess Entry. *American Economic Review*, 89(1), 306—318.
- De Paola, M., Scoppa, V., & Lombardo, R. (2010). Can Gender Quotas Break Down Negative Stereotypes? Evidence from Changes in Electoral Rules. *Journal of Public Economics*, 94(5/6), 344—353.
- Ellsberg, D. (1961). Risk, Ambiguity and Savage Axioms. *Quarterly Journal of Economics*, 75(4), 643—679.

- Ferber, M. & Waldfogel, J. (1998). The Long-term Consequences of Non-standard Work. *Monthly Labor Review*, 121(5), 3–12.
- Ferrari, S. & Cribari-Neto, F. (2004). Beta Regression for Modelling Rates and Proportions. *Journal of Applied Statistics*, 31(7), 799–815.
- Fraley, C. & Raftery, A. (2007). Bayesian Regularization for Normal Mixture Estimation and Model-Based Clustering. *Journal of Classification*, 24(2), 155–181.
- Gans, H., Knox, G., & Croson, R. (2007). Simple Models of Discrete Choice and Their Performance in Bandit Experiments. *Manufacturing & Service Operations Management*, 9(4), 383–408.
- GEM (2015). gemconsortium.org.
- Gittins, J., Glazebrook, K., & Weber, R. (2011). *Multi-armed Bandit Allocation Indices* (2. ed.). Chichester: Wiley.
- Gittins, J. & Jones, D. (1974). A Dynamic Allocation Index for the Sequential Design of Experiments. In J. Gani (Ed.), *Progress in Statistics* (pp. 241–266). Amsterdam, NL: North-Holland.
- Granrose, C. & Baccili, P. (2006). Do Psychological Contracts Include Boundaryless or Protean Careers? *Career Development International*, 11(2), 163–182.
- Griffin, D. & Varey, C. (1996). Towards a Consensus on Overconfidence. *Organizational Behavior and Human Decision Processes*, 65(3), 227–231.
- Hall, D. & Moss, J. (1998). The New Protean Career Contract: Helping Organizations and Employees Adapt. *Organizational Dynamics*, 26(3), 22–37.
- Hamilton, B. (2000). Does Entrepreneurship Pay? *Journal of Political Economy*, 108(3), 604–631.
- Hartigan, J. & Wong, M. (1979). A K-Means Clustering Algorithm. *Journal of the Royal Statistical Society*, 28(1), 100–108.
- Hofstede, G. (2015). www.geerthofstede.nl/research--vsm.
- Hofstede, G., Hofstede, G., & Minkov, M. (2010). *Cultures and Organizations* (3. ed.). New York: McGraw-Hill.
- Holmes, T. & Schmitz, J. (1990). A Theory of Entrepreneurship and Its Application to the Study of Business Transfers. *Journal of Political Economy*, 98(2), 265–294.
- Huang, R. (2007). Distance and Trade: Disentangling Unfamiliarity Effects and Transport Cost Effects. *European Economic Review*, 51(1), 161–181.
- Huang, R. (2008). Tolerance for Uncertainty and the Growth of Informationally Opaque Industries. *Journal of Development Economics*, 87(2), 333–353.
- Hundley, G. (2001). Why and When are the Self-employed More Satisfied with Their Work? *Industrial Relations*, 40(2), 293–316.

- Inklaar, R. & Yang, J. (2012). The Impact of Financial Crises and Tolerance for Uncertainty. *Journal of Development Economics*, 97(2), 466–480.
- John, O. & Srivastava, S. (2001). The Big-Five Trait Taxonomy. In L. Pervin & O. John (Eds.), *Handbook of Personality* (2. ed.). chapter 4, (pp. 102–138). New York: Guilford Press.
- Jovanovic, B. (1979). Job Search and the Theory of Turnover. *Journal of Political Economy*, 87(5), 972–990.
- Kelley, D., Brush, C., Greene, P., & Litovsky, Y. (2013). *Global Entrepreneurship Monitor: 2012 Women's Report*. London: London Business School.
- Kihlstrom, R. & Laffont, J.-J. (1979). A General Equilibrium Entrepreneurial Theory of Firm Formation Based on Risk Aversion. *Journal of Political Economy*, 87(4), 719–748.
- Koellinger, P., Minniti, M., & Schade, C. (2007). I Think I Can, I Think I Can. *Journal of Economic Psychology*, 28(4), 502–527.
- Kosmidis, I. & Firth, D. (2010). A Generic Algorithm for Reducing Bias in Parametric Estimation. *Electronic Journal of Statistics*, 4, 1097–1112.
- Lazear, E. (2005). Entrepreneurship. *Journal of Labor Economics*, 23(4), 649–680.
- Lucas, R. (1978). On the Size Distribution of Business Firms. *Bell Journal of Economics*, 9(2), 508–523.
- MacDonald, G. (1988). The Economics of Rising Stars. *American Economic Review*, 78(1), 155–166.
- McCrae, R. & John, O. (1992). An Introduction to the Five-Factor Model and Its Applications. *Journal of Personality*, 60(2), 175–215.
- Meyer, R. & Shi, Y. (1995). Sequential Choice under Ambiguity. *Management Science*, 41(5), 817–834.
- OECD (2015). *Entrepreneurship at a Glance 2015*. Paris: OECD Publishing.
- Ramírez, A. & Tadesse, S. (2009). Corporate Cash Holdings, Uncertainty Avoidance, and the Multinationality of Firms. *International Business Review*, 18(4), 387–403.
- Rauch, A. & Frese, M. (2007). Let's Put the Person back into Entrepreneurship Research. *European Journal of Work and Organizational Psychology*, 16(4), 353–385.
- Roessler, C. & Koellinger, P. (2012). Entrepreneurship and Organization Design. *European Economic Review*, 56(4), 888–902.
- Rothschild, M. (1974). A Two-Armed Bandit Theory of Market Pricing. *Journal of Economic Theory*, 9(2), 185–202.
- Savage, L. (1954). *The Foundations of Statistics*. New York: Wiley.
- Sherman, R. (1974). The Psychological Difference Between Ambiguity and Risk. *Quarterly Journal of Economics*, 88(1), 166–169.

- Steyvers, M., Lee, M., & Wagenmakers, E.-J. (2009). A Bayesian Analysis of Human Decision-making on Bandit Problems. *Journal of Mathematical Psychology*, 53(3), 168–179.
- Stuetzer, M., Goethner, M., & Cantner, U. (2012). Do Balanced Skills Help Nascent Entrepreneurs to Make Progress in the Venture Creation Process? *Economics Letters*, 117(1), 186 – 188.
- Sullivan, J. & Peterson, R. (1991). A Test of Theories Underlying the Japanese Lifetime Employment System. *Journal of International Business Studies*, 22(1), 79–97.
- Sullivan, S. & Baruch, Y. (2009). Advances in Career Theory and Research: A Critical Review and Agenda for Future Exploration. *Journal of Management*, 35(6), 1542–1571.
- Tobin, J. (1958). Liquidity Preference as Behavior Towards Risk. *The Review of Economic Studies*, 25(2), 65–86.
- Valcour, P. & Tolbert, P. (2003). Gender, Family and Career in the Era of Boundarylessness: Determinants and Effects of Intra- and Inter-organizational Mobility. *International Journal of Human Resource Management*, 14(5), 768–787.
- Verbruggen, M., van Emmerik, H., Van Gils, A., Meng, C., & de Grip, A. (2015). Does Early-career Underemployment Impact Future Career Success? A Path Dependency Perspective. *Journal of Vocational Behavior*, 90(1), 101–110.
- Wagner, J. (2006). Are Nascent Entrepreneurs ‘Jacks-Of-All-Trades’? A Test Of Lazear’s Theory of Entrepreneurship with German Data. *Applied Economics*, 38(20), 2415–2419.
- Weber, R. (1992). On the Gittins Index for Multiarmed Bandits. *The Annals of Applied Probability*, 2(4), 1024–1033.
- Williams, D. (2000). Consequences of Self-employment for Women and Men in the United States. *Labour Economics*, 7(5), 665–687.
- World Bank (2015). <http://data.worldbank.org/>.
- Zhao, H. & Seibert, S. (2006). The Big Five Personality Dimensions and Entrepreneurial Status. *Journal of Applied Psychology*, 91(2), 259–271.

Appendix A

This appendix demonstrates the properties of the behavioral rule mentioned in the main text.

Proof of (PB1). Let $b_{i,t} := b(x_{i,t})$. Furthermore, let $b_{j,t} = \max\{b_{w,t}, b_{e,t}\}$. Assume that in period t we are required to select a safe option u over a non-safe option $j \neq u$ because $b_{j,t-1} \leq 0$. Then, $b_{i,t} = b_{i,t-1}$ for all $i \in \mathbb{O}$ because we do not obtain any new information by selecting unemployment. Consequently, $b_{j,t} \leq 0$. Hence: If $b_{j,t} \leq 0$, then $b_{j,h} \leq 0$ for all $h \in \{t+1, t+2, \dots\}$ and $j \neq u$ by induction. There are two interesting starting points. First, we will start with u in period 1 if $b_{j,0} \leq 0$ and there will be no transitions at all. Second, we will start with option $j \neq u$ if $b_{j,0} > 0$ and there will be no changes of occupational status after t^* if for some $t^* > 1$ we have $b_{j,t^*-1} \leq 0$. ■

Proof of (PB2). We show the existence of unusual individual types: Type- G individuals gamble for the sake of gambling, while type- P individuals are paralyzed by uncertainty.

TYPE G. An option generating a success with a known probability of 100% has the behavioral index $b = 1$. Thus, there exists a critical value $\lambda_G(x, \gamma_b) > 0$ such that $b(x; \lambda, \gamma_b) > 1$ if $\lambda > \lambda_G$ where

$$\lambda_G(x, \gamma_b) = \frac{\beta}{\gamma_b \left[\frac{\alpha\beta}{\beta^3 + (3\alpha+1)\beta^2 + (3\alpha^2+2\alpha)\beta + \alpha^3 + \alpha^2} \right]^{\frac{1}{2}}} > 0$$

If $\lambda > \lambda_G(x_t, \gamma_b)$ for some t , the gambler will gamble for the sake of gambling for at least one period. Conditions for being of type G are quite restrictive as they demand (a high level of) uncertainty affinity.

TYPE P. The behavioral index of unemployment is zero. An individual is paralyzed if all options $i \in \mathbb{O}_{-u}$ are uncertain and $\lambda \leq \min\{\lambda_P(x_e, \gamma_b), \lambda_P(x_w, \gamma_b)\}$ such that $\max\{b(x_e; \lambda, \gamma_b), b(x_w; \lambda, \gamma_b)\} \leq 0$. The critical value is

$$\lambda_P(x, \gamma_b) = -\frac{\alpha}{\gamma_b \left[\frac{\alpha\beta}{\beta^3 + (3\alpha+1)\beta^2 + (3\alpha^2+2\alpha)\beta + \alpha^3 + \alpha^2} \right]^{\frac{1}{2}}} < 0$$

Note that if $\lambda \leq \min\{\lambda_P(x_{e,0}, \gamma_b), \lambda_P(x_{w,0}, \gamma_b)\}$, no other option than unemployment will ever be selected. Paralysis requires relatively modest levels of uncertainty aversion if the number of observations is low. Hence, being paralyzed is more realistic than being a gambler for the sake of gambling. ■

Proof of (PB3). The behavioral rule does not always stay on a winner. A success can make an option more uncertain if not expected and an uncertainty-averse gambler might switch from a winner. Furthermore, there are individuals who like uncertainty and who would switch from a successful option because a success makes an option less uncertain. Let i be an option producing a success and j a second option. Assume

that

$$b(\alpha_i, \beta_i; \lambda, \gamma_b) > b(\alpha_j, \beta_j; \lambda, \gamma_b)$$

such that option i is selected. After option i is selected, it is assumed to generate a success. Consequently, if $b(\alpha_i + 1, \beta_i; \lambda, \gamma_b) \geq b(\alpha_i, \beta_i; \lambda, \gamma_b)$, the gambler would always stay on a winner. If $b(\alpha_i + 1, \beta_i; \lambda, \gamma_b) < b(\alpha_i, \beta_i; \lambda, \gamma_b)$, the gambler does not necessarily stay on a winner since we might have $b(\alpha_i + 1, \beta_i; \lambda, \gamma_b) < b(\alpha_j, \beta_j; \lambda, \gamma_b)$. Put differently, if we have

$$\lambda \gamma_b z(\alpha, \beta) > \mu(\alpha + 1, \beta) - \mu(\alpha, \beta), \quad z(\alpha, \beta) := \frac{\sigma(\alpha, \beta)}{\gamma} - \frac{\sigma(\alpha + 1, \beta)}{1 + \gamma}$$

the gambler might not stay on a winner. Note that $\mu(\alpha + 1, \beta) - \mu(\alpha, \beta) > 0$. In most cases, we have $z(\alpha, \beta) > 0$ and not staying on a winner requires

$$\lambda > \frac{\mu(\alpha + 1, \beta) - \mu(\alpha, \beta)}{\gamma_b z(\alpha, \beta)} > 0$$

such that uncertainty-averse individuals will always stay on a winner. However, if a success was not expected, we could have $z(\alpha, \beta) < 0$. For instance, if $\alpha = 1$ and $\beta = 10$, which is an event with a positive probability if $\phi \in (0, 1)$, a success is not expected such that $z(1, 10) < 0$. In such a case, the gambler will not stay on a winner and switch from option i to j if uncertainty aversion is sufficiently high or, more precisely, if

$$\lambda < \frac{\mu(\alpha_j, \beta_j) - \mu(\alpha_i + 1, \beta_i)}{\gamma_b \left[\sigma(\alpha_i + 1, \beta_i) [\gamma_i + 1]^{-1} - \sigma(\alpha_j, \beta_j) \gamma_j^{-1} \right]}$$

■

Proof of (PB4). The property is easy to verify if we consider a choice between unemployment and a non-unemployment option. We can always find a $\lambda \in \mathbb{R}^-$ such that $b_{i,h} \leq 0$ for all $i \in \mathbb{O}_{-u}$ and some period $h > 0$. Given that unemployment is selected in one period, unemployment is selected in all consecutive periods (see property [PB1]). This is independent from true success probabilities ϕ_e and ϕ_w such that the best option might not be identified. However, the property also holds if we assume that unemployment is never weakly better than a non-unemployment option:

$$q_u := \mathbb{P} \{ b_{i,t} > 0 \text{ for all } i \in \mathbb{O}_{-u} \text{ and all } t \} = 1$$

Note that $q_u = 1$ requires

$$\lambda > -\frac{\gamma_{i,t} \mu_{i,t}}{\gamma_b \sigma_{i,t}} \quad \text{for all } i \in \mathbb{O}_{-u} \text{ and all } t$$

which, for instance, holds if individuals are uncertainty-affine. The following argument is a modified version of Theorem 2 in Brezzi & Lai (2000). To see that the property holds in a situation where $q_u = 1$, first, consider the following two condi-

tions:

$$\mathbb{P} \left\{ \lim_{t \rightarrow \infty} \mu_{i,t} = \phi_i \right\} = 1 \quad (\text{A.1})$$

$$\mathbb{P} \left\{ \lim_{t \rightarrow \infty} \hat{\sigma}_{i,t} = 0 \right\} = 1 \quad (\text{A.2})$$

where $\hat{\sigma}_{i,t} := \sigma_{i,t}/\gamma_{i,t}$. Note that t is, here, the number of times option i is sampled from. To get (A.1), use the assumption that the initial prior $\mathbf{x}_0 = (\alpha_0, \beta_0)$ is sampled from a distribution with the right success probability ϕ , and note that

$$\mu_t = \frac{\alpha_t}{\gamma_t} = \frac{1}{t} B_t, \quad B_t := \sum_{h=1}^t b_h, \quad b_t \sim \text{Bernoulli}(\phi), \quad \mathbb{E}[b] = \phi$$

where I slightly abuse notation to extend the number of observed periods to pseudo observations, \mathbf{x}_0 . By applying the Strong Law of Large Numbers, we obtain (A.1). To show (A.2), note that if pseudo observations are sampled from the correct distribution, we have $\alpha_t = B_t$ and $\beta_t = t - B_t$, such that

$$\hat{\sigma}_t = \frac{1}{t^2} \left[\frac{B_t(t - B_t)}{t^3 + t^2} \right]^{\frac{1}{2}}$$

where $0 \leq B_t \leq t$. Now, notice that for a fixed t the expression $B_t(t - B_t)$ takes a maximum at $B_t = t/2$ such that we get

$$\hat{\sigma}_t = \frac{1}{t^2} \left[\frac{B_t(t - B_t)}{t^3 + t^2} \right]^{\frac{1}{2}} \leq \frac{1}{2t[t^2(t + 1)]^{\frac{1}{2}}} \quad \text{for all } t$$

By letting $t \rightarrow \infty$, we obtain (A.2). If (A.1) and (A.2) are combined, we get

$$\mathbb{P} \left\{ \lim_{t \rightarrow \infty} b_{i,t} = \lim_{t \rightarrow \infty} \mu_{i,t} + \lambda \gamma_b \lim_{t \rightarrow \infty} \hat{\sigma}_{i,t} = \phi_i \right\} = 1 \quad (\text{A.3})$$

According to (A.3), an infinite amount of information is equivalent to perfect information. Hence, if the gambler selects an occupational option for an infinite amount of periods, she will almost surely learn the true success probability of this occupational option. Yet, only one occupation will be selected for an infinite number of periods. To see this, define the number of periods option $i \in \mathbb{O}_{-u}$ is selected by $T_{i,t}$ and let $T_i^\infty := \lim_{t \rightarrow \infty} T_{i,t}$. On the event $T_e^\infty = \infty$, $b_{e,t} \rightarrow \phi_e$ almost surely. On the event $T_w^\infty = \infty$, $b_{w,t} \rightarrow \phi_w$ almost surely. But, since only the occupation with the highest index is selected, we cannot simultaneously have $T_e^\infty = \infty$ and $T_w^\infty = \infty$ if $\phi_e \neq \phi_w$ or

$$\mathbb{P} \{ T_i^\infty < \infty \text{ for all } i \in \mathbb{O}_{-u} \text{ except one } i \} = 1$$

It can be demonstrated that the best option i^* is not necessarily the option selected for an infinite number of periods or

$$\mathbb{P} \{ T_{i^*}^\infty < \infty \} > 0 \quad (\text{A.4})$$

To obtain (A.4), assume, without loss of generality, $i^* = w$ such that $\phi_w > \phi_e$. Let

$$q_e := \mathbb{P} \{b_{e,t} \geq \varepsilon \text{ for all } t > 0\} \quad (\text{A.5})$$

for some ε and let

$$q_w := \mathbb{P} \{b_{w,k} < \varepsilon\} \quad (\text{A.6})$$

for some k . Consider the joint probability of the event in (A.5) and (A.6) given by

$$q^* := \mathbb{P} \{b_{w,k} < \varepsilon \leq b_{e,t} \text{ for all } t > 0\}$$

If $q^* > 0$, (A.4) holds. Note that b_e and b_w are independent as the underlying options we sample from are independent. By independence, we get $q^* = q_e q_w$. Let $\varepsilon \in (0, \phi_e)$. Given $q_u = 1$ and (A.3), we can always find a ε such that $q_e > 0$. It is also obvious that $q_w > 0$. For instance: Let $x_{w,0} = (1, 1)$, which occurs with probability $\phi_w[1 - \phi_w] > 0$ if the initial prior is sampled from a distribution with success probability ϕ_w . Assume that we obtain k nonsuccesses. The probability to obtain k nonsuccesses is given by $[1 - \phi_w]^k > 0$. In such a setting, we get

$$b_{w,k} = \frac{\gamma_b \left[\frac{k+1}{k^3+7k^2+16k+12} \right]^{\frac{1}{2}} \lambda + 1}{k+2}$$

such that

$$\frac{db_{w,k}}{dk} = -\frac{\sqrt{k+3}\gamma_b(2k^2+7k+4)\lambda + \sqrt{k+1}(k^3+8k^2+21k+18)}{\sqrt{k+1}(k^5+12k^4+57k^3+134k^2+156k+72)} < 0$$

if

$$\lambda > -\frac{\sqrt{k+1}(k^3+8k^2+21k+18)}{\sqrt{k+3}(2k^2+7k+4)\gamma_b}$$

As $b_{w,k}$ monotonically approaches zero as k increases, it is always possible to generate a sequence of events with a positive probability such that $b_{w,k} < \varepsilon$. Hence: $q^* > 0$. ■

Appendix B

This appendix compares choice patterns produced by the behavioral rule to patterns obtained in Bernoulli bandit experiments. Most experiments analyze the explanatory performance of the optimal decision rule and try to classify deviations from optimal behavior.

Deviations are quite common (see Meyer & Shi 1995; Anderson 2001; Gans et al. 2007). Some results are strongly consistent with the behavioral rule. For instance, Anderson (2001) finds that individuals experiment less than optimal and that this phenomenon is best explained by uncertainty aversion. According to Anderson (2001), individuals follow a rule with a lower than the optimal index. Brezzi & Lai

(2000) demonstrated that the optimal index has the following boundaries: $\mu \leq g \leq \mu + \delta(1 - \delta)^{-1}\sigma$. An uncertainty-averse gambler deciding according to the behavioral rule has an index $b < \mu \leq g$. Put differently, an uncertainty-averse gambler following the behavioral rule operates with an index that is always lower than the optimal Gittins index.

Two empirical findings that are not obviously consistent with the behavioral rule are as follows.

- (R1) Steyvers et al. (2009) calculate behavioral characteristics (denoted henceforth by BCHAR) for individuals experimenting with Bernoulli bandits with varying success distributions. A BCHAR is, for instance, the number of times individuals selected an option with fewer successes and nonsuccesses than an alternative option, which is interpreted as a measure of exploratory behavior. Steyvers et al. (2009) find no significant correlation between any of the BCHARS and personality traits.
- (R2) Gans et al. (2007) apply different rules to data generated by individuals who played with Bernoulli bandits. Gans et al. (2007) establish a ranking of the explanatory performance of different rules.

The first result (R1) appears to contradict my assumption that decisions are affected by personality (through the parameter λ). The ranking of Gans et al. (2007), result (R2), should be reproducible with data generated by the behavioral rule. I show that both empirical findings are consistent with the behavioral rule.

B.1 Personality and simple characteristics of behavior

The rationale behind the first test is based on the idea that personality can determine individual levels of uncertainty aversion but aspects of uncertainty-averse behavior can look exactly the same across different levels of aversion. In other words, uncertainty aversion does not have a linear effect on behavior. A BCHAR might capture the difference between an uncertainty-averse and an uncertainty-neutral individual but BCHARS may not capture the difference between two uncertainty-averse individuals.

Let $\chi \in \mathbb{R}$ denote an observable personality trait. Let $\mathbf{d}_{\chi, \lambda}$ denote a draw from the joint distribution of the observable personality trait and $\lambda \in \mathbb{R}$ representing uncertainty preferences. The joint distribution is bivariate normal. I assume that personality almost completely determines the parameter λ such that the correlation between χ and λ is 0.99 (almost perfect). The personality trait has a variance of 1; the uncertainty-preferences parameter has a variance of 2; and the mean of the personality trait is zero.

There are four groups of individuals. Each group corresponds to 100 draws of $\mathbf{d}_{\chi, \lambda}$.

FIRST GROUP: For this group, I assume that the mean of λ is -50 . Individuals from group 1 are uncertainty averse with a probability of almost 100%.

SECOND GROUP: For this group, the mean of λ is -10 . Individuals from group 2 are uncertainty averse with a very high probability but the probability is slightly smaller than for individuals from group 1.

THIRD GROUP: The mean of λ is zero. A member of group 3 is uncertainty averse with a probability of 50%.

FOURTH GROUP: For the last group, I assume that the mean of λ is 10 such that group-4 individuals are uncertainty averse with a probability of around 0%.

The Monte Carlo experiment proceeds as follows. First, groups are drawn. Second, individuals in each group use the behavioral rule to play with a two-armed bandit over 100 periods. The bandit setting corresponds to the setting used by Steyvers et al. (2009). After each individual finishes the game, five BCHARS, which are described below, are calculated. Then, the algorithm computes the correlation between BCHARS and the observable personality trait. This procedure is repeated 1,000 times for each combination of priors and success probabilities.

The five BCHARS, which are also calculated by Steyvers et al. (2009), are as follows. **EXPLOITATION** is the number of times the gambler selected an option with more successes and more nonsuccesses than the alternative option. **EXPLORATION** is the number of times the gambler selected an option with fewer successes and fewer nonsuccesses than the alternative option. **BETTER** is the number of times the gambler selected an option with more success and fewer nonsuccesses than the alternative option. **WORSE** is the number of times the gambler selected an option with fewer successes and more nonsuccesses than the alternative option. **UNTRIED** is the number of times the gambler selected an option never selected before.

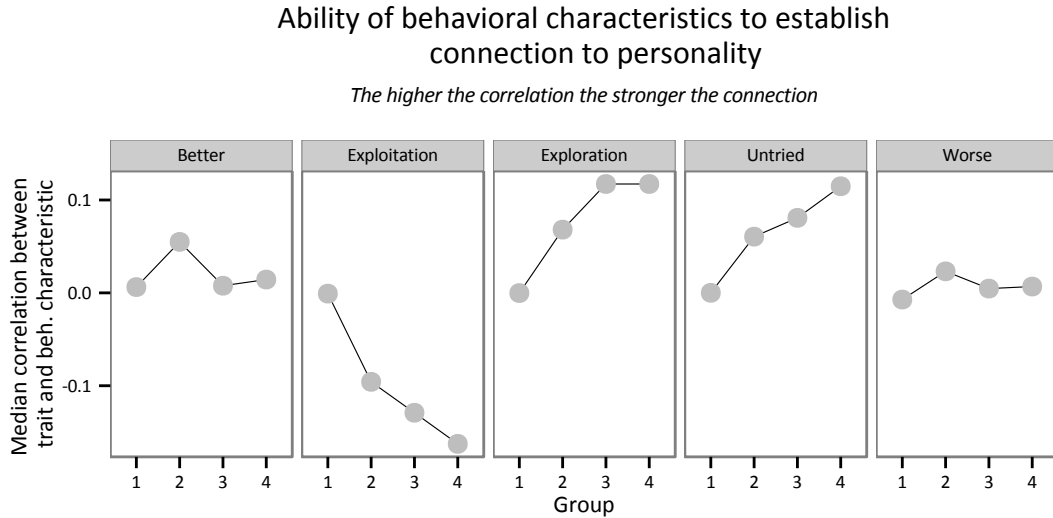
Decisions can be influenced by decision noise. Noise is defined as follows.

DEFINITION B. (Decision noise) A rule-based behavior is said to be affected by decision noise $\epsilon \in (0, 0.5]$ if, in some situation, given two different λ options 1 and 2, a decision rule prescribes to select (a) option 1 with probability 1 and option 2 with probability zero or (b) option 1 with probability zero and option 2 with probability 1 but the gambler selects (a) option 1 with probability $1 - \epsilon$ and option 2 with probability ϵ or (b) option 1 with probability ϵ and option 2 with probability $1 - \epsilon$.

ϵ is the probability to make a mistake by deviating from the prescribed choice probability. I set $\epsilon = 0.01$ such that there are only very few mistakes.

Figure B.1 presents median correlations between the personality trait and BCHARS conditional on group affiliation. **BETTER** and **WORSE** do not capture any trait differences inside groups. **EXPLOITATION**, **EXPLORATION**, and **UNTRIED** reflect trait differences inside groups if there are some uncertainty-neutral or uncertainty-seeking individuals in the group. The correlations are plausible. For instance, an increase in the personality trait, which goes along with a reduction in uncertainty aversion, is associated with a decrease in **EXPLOITATION**-behavior and an increase in **EXPLORATION**-behavior. However, given that there is an almost perfect correlation between the personality trait and the uncertainty-aversion parameter λ determining behavior, correlations between the personality trait and BCHARS are rather small. The correlations also tend to decrease with the share of uncertainty-averse individuals in groups.

FIGURE B.1: BCHARS and personality



For instance, in the first group, which is completely uncertainty-averse, no BCHAR can establish a meaningful connection between personality and behavior—every individual in the totally uncertainty-averse group, who does not make a mistake, exploits but not explores and never tries an untried option (besides the first period). Since a large share of uncertainty-averse individuals (in an arbitrary group) is a realistic scenario, there is a simple explanation for the results of Steyvers et al. (2009).

B.2 Explanatory performance

Gans et al. (2007) use a Bernoulli bandit experiment to estimate the explanatory performance of different decision rules. Among others, Gans et al. (2007) consider the following three rules: the myopic rule, the simple rule, and exponential smoothing. Explanatory performance is measured by the Bayesian information criterion (BIC). According to Gans et al. (2007), exponential smoothing performs better than the simple rule, while the simple rule performs better than the myopic rule. In what follows, I will describe the three rules and show that applying the rules to data generated by the behavioral rule results in the ranking of Gans et al. (2007).

The rules are interesting because they are derived from different assumptions. Alternative rules are as follows.

MYOPIC RULE. The myopic rule, r^M , is a special case of the Gittins index with a zero probability to continue. In the context of the behavioral rule, myopia corresponds to the behavior of uncertainty-neutral individuals. The myopic index is

$$r_{i,t}^M = \mu_{i,t} \quad (\text{B.1})$$

SIMPLE RULE. According to the simple rule, r^S , an option can have two qualities: It is either good or bad but never both. The index of the simple rule is (see Gans et al.

2007):

$$r_{i,t}^S = r_{i,t-1}^S + (a_0 + a)\pi_t - a \quad (\text{B.2})$$

which corresponds to a random walk. Gans et al. (2007) set $a_0 = 1$ and $a > 0$ is a free parameter.

EXPONENTIAL SMOOTHING RULE. Exponential smoothing, r^{ES} , is a purely descriptive rule. The smoothing index is a weighted average of the smoothing index from the previous period and new information:

$$r_{i,t}^{\text{ES}} = (1 - \zeta)r_{i,t-1}^{\text{ES}} + \zeta\pi_t \quad (\text{B.3})$$

where $\zeta \in [0, 1]$ is a weight. The weight is a free parameter. In addition, starting values, i.e., $r_{i,0}^{\text{ES}}$, are also free parameters. In line with Gans et al. (2007), I assume that the initial index of one option is $2/3$.

The Monte Carlo experiment proceeds as follows. First, the algorithm generates a decision history based on the behavioral rule and some combination of uncertainty preferences, probabilities to succeed, and noise levels. Decision histories are generated for 100 periods, while the choice setting is adopted from Gans et al. (2007). There are two noise levels: $\epsilon = 0.03$ is interpreted as low noise, whereas $\epsilon = 0.08$ is interpreted as high noise. Then, BIC values are estimated by maximum likelihood with a probit model. Each model has a sensitivity parameter. The myopic model has no additional free parameters. The model based on the simple rule has one additional parameter. The exponential smoothing model has two additional parameters. The two steps are sequentially repeated 1,000 times.

Figure B.2 presents selected results. Note that, here, a smaller BIC value indicates higher explanatory performance. If we consider a setting with low noise and sufficiently high uncertainty aversion ($\lambda \leq -4$), we can observe the ranking established by Gans et al. (2007), i.e., exponential smoothing is better than the simple rule, while the simple rule is better than the myopic rule. Given high noise, the appropriate ranking is produced for high uncertainty aversion ($\lambda = -20$) only. Therefore, there exist parameter combinations where the behavioral rule is rank-consistent with empirical evidence. Pronounced uncertainty aversion and a low level of decision noise characterize these combinations. Put differently, if most decisions can be derived from the behavioral rule with uncertainty aversion and are not random, theoretical behavior is similar to the behavior of actual gamblers.

Appendix C

This appendix provides data characteristics. The data set features three important variables: uncertainty preferences, approximated by Hofstede's uncertainty avoidance index; average established business activity; and enterprise death rates. Furthermore, I use additional covariates (viz., GDP growth, inflation, real interest rate, start-up costs, unemployment rates, and average ease of doing business).

Countries in data sets used to assess the two predictions are shown in Table C.1.

FIGURE B.2: Performance rankings with low and high decision noise

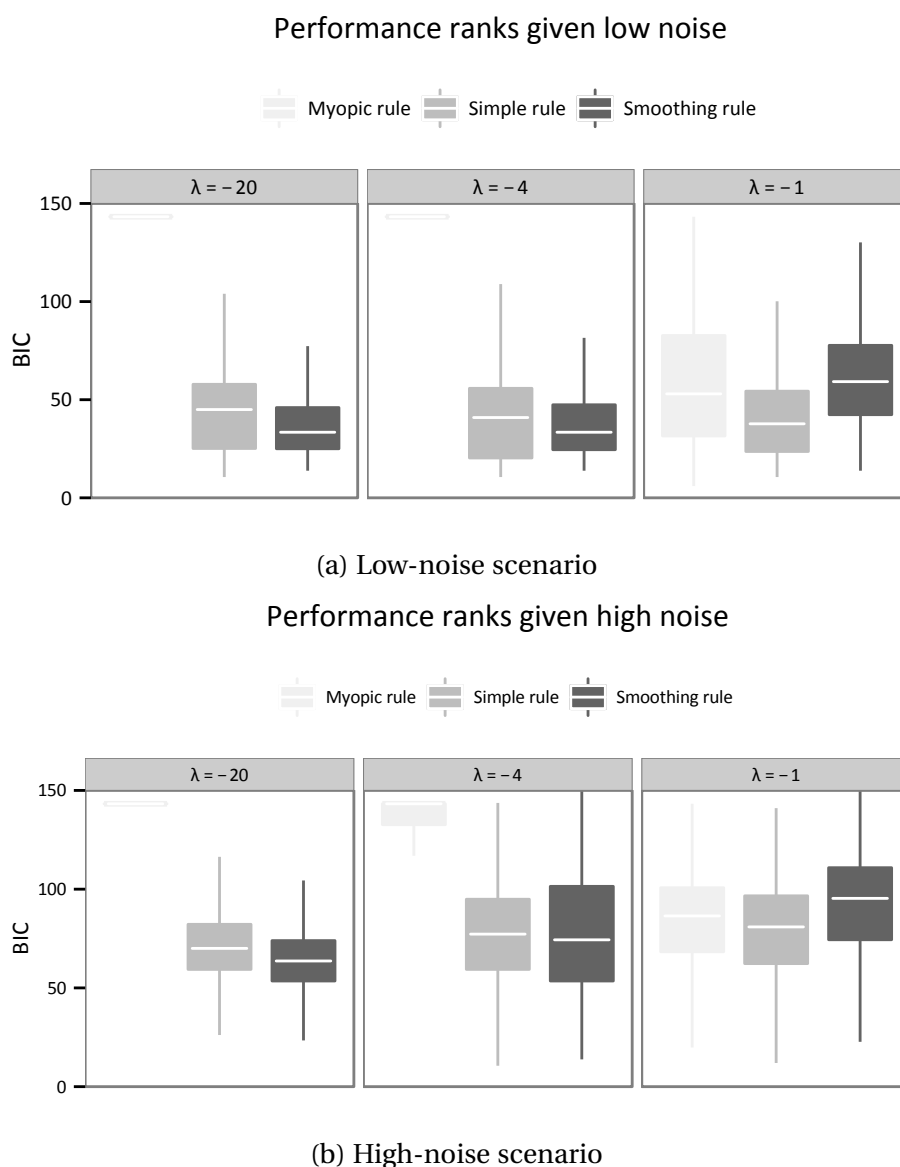


Table C.2 presents a detailed analysis of Hofstede's index. The table is the result of the application of k -means clustering with three clusters, where I used the Hartigan & Wong (1979) algorithm to minimize within-cluster sums of squares. According to Table C.2, cluster means explain ca. 87% of the overall variation (100% is perfect fit). The high-aversion group consists of approx. 48% of all countries with an available index. The moderate- and high-aversion group accounts for ca. 84% of all countries, while the low-aversion group accounts only for 16%. To summarize, most societies have either a moderate or a high level of uncertainty aversion.


The Global Entrepreneurship Monitor defines the established business ownership rate as the percentage of the population aged 18–64 who run or manage a business that paid payments to the owner for more than 42 months. Figure C.1 shows the distribution of *average* established business activity. The distribution is skewed such

TABLE C.1: Countries in data

FIRST PREDICTION	SECOND PREDICTION
Argentina, Australia, Austria, Bangladesh, Belgium, Brazil, Canada, Chile, China, Colombia, Costa Rica, Croatia, Czech Rep., Denmark, Ecuador, El Salvador, Estonia, Finland, France, Germany, Great Britain, Greece, Guatemala, Hong Kong, Hungary, India, Indonesia, Iran, Ireland, Israel, Italy, Jamaica, Japan, Korea South, Latvia, Lithuania, Luxembourg, Malaysia, Mexico, Morocco, Netherlands, New Zealand, Norway, Pakistan, Panama, Peru, Philippines, Poland, Portugal, Romania, Russia, Serbia, Singapore, Slovak Rep., Slovenia, Spain, Suriname, Sweden, Switzerland, Taiwan, Thailand, Trinidad and Tobago, Turkey, USA, Uruguay, Venezuela, Vietnam	Australia, Belgium, Brazil, Canada, Czech Rep., Estonia, Hungary, Israel, Italy, Korea South, Latvia, Lithuania, Netherlands, New Zealand, Romania, Slovak Rep., Slovenia, USA
67 countries	18 countries

that low values are more likely than high values. Consequently, most countries have a rather low average rate of established business activity.

FIGURE C.1: Distribution of average established business activity rates (white line is median)



OECD (2015) defines firm death as follows.

An employer enterprise death occurs either at the death of an enterprise with at least one employee in the year of death or when an enterprise shrinks to below the threshold of one employee for at least two years. ... The employer enterprise death rate corresponds to the number of deaths of employer enterprises as a percentage of the population of active enterprises with at least one employee. (OECD 2015, p. 52)

TABLE C.2: Low-, moderate-, and high-aversion groups

	LOW-AVERSION GROUP	MODERATE- AVERSION GROUP	HIGH-AVERSION GROUP
GROUP MEAN OF HOFSTEDE'S INDEX	28.00	58.48	88.09
COUNTRIES BELONGING TO GROUP	China, Denmark, Great Britain, Hong Kong, India, Ireland, Jamaica, Malaysia, Singapore, Sweden, Vietnam	Australia, Austria, Bangladesh, Canada, Ecuador, Estonia, Finland, Germany, Indonesia, Iran, Latvia, Lithuania, Luxembourg, Morocco, Netherlands, New Zealand, Norway, Pakistan, Philippines, Slovak Rep., Switzerland, Taiwan, Thailand, Trinidad and Tobago, USA	Argentina, Belgium, Brazil, Bulgaria, Chile, Colombia, Costa Rica, Croatia, Czech Rep., El Salvador, France, Greece, Guatemala, Hungary, Israel, Italy, Japan, Korea South, Malta, Mexico, Panama, Peru, Poland, Portugal, Romania, Russia, Serbia, Slovenia, Spain, Suriname, Turkey, Uruguay, Venezuela
GROUP SIZE	11	25	33
WITHIN SUM OF SQUARES	966.00	1714.24	2360.73
TOTAL VARIATION ACCOUNTED FOR BY GROUP MEANS			86.8%

FIGURE C.2: Distribution of firm death rates (white line is median)

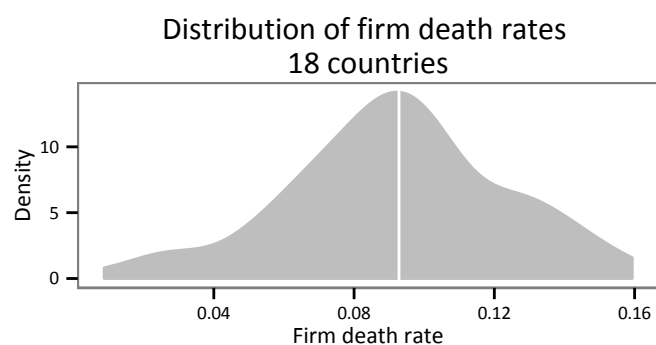
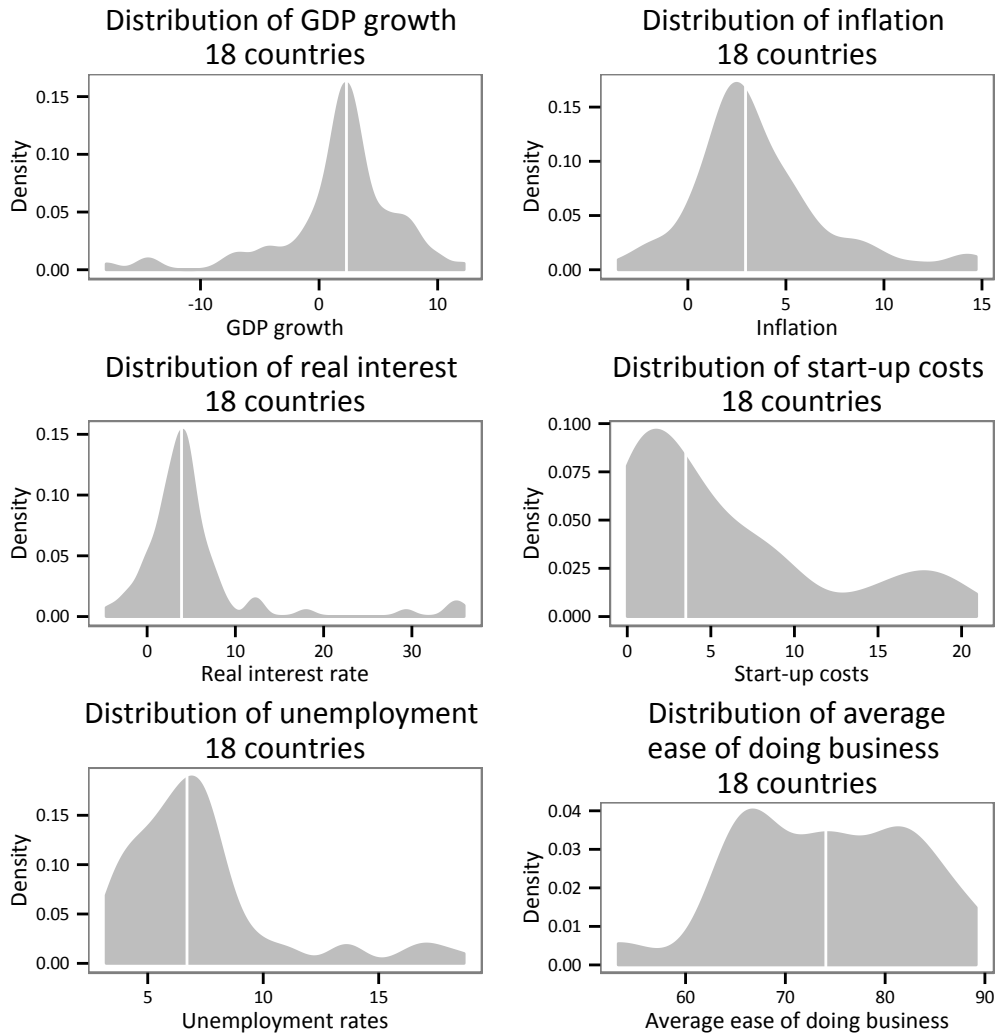


Figure C.2 plots the distribution of firm death rates, which appears fairly normal such that it is not surprising that OLS and the beta regression deliver similar answers.

FIGURE C.3: Distributions of GDP growth, inflation, real interest rates, start-up costs, unemployment rates, and average ease of doing business (white vertical lines represent medians)



In Figure C.3, I plot the distributions of GDP growth, inflation, real interest rates, start-up costs, unemployment rates, and average ease of doing business.