The Journal of Experimental Biology 210, 269-277
Published by The Company of Biologists 2007
doi:10.1242/jeb.02656

Parameters of variable reward distributions that affect risk sensitivity of honey bees

Tamar Drezner-Levy and Sharoni Shafir*

B. Triwaks Bee Research Center, Department of Entomology, Faculty of Agricultural, Food and Environmental Quality Sciences, The Hebrew University of Jerusalem, Rehovot, 76100, Israel

*Author for correspondence (e-mail: shafir@agri.huji.ac.il)

Accepted 14 November 2006

Summary

We investigated risk sensitivity with harnessed honey bees in a proboscis-extension conditioning paradigm. We conditioned each subject to turn its head and extend its proboscis towards one of two presented odors; one odor was associated with a constant reward and the other with a variable reward that was either low or high, with probabilities *P* and (1–*P*), respectively. Reward values and probabilities were set so that the expected value of the variable alternative was equal to that of the constant one. We performed six experimental conditions in which variability was in reward volume and three conditions in which variability was in reward concentration. The experiments were designed to systematically test the effect of various parameters that describe the reward distributions on levels of risk sensitivity. Risk aversion was greatest when variability was in reward volume, and the variable distribution included zero rewards and had a high coefficient of variation (CV=s.d./mean). The variance itself did not affect risk sensitivity. Subjects were risk indifferent when the variable distribution did not include zero rewards, however these distributions were positively skewed. The independent effects of zero rewards and distribution skew remain to be tested. Subjects were risk indifferent in conditions where variability was in reward concentration, but concentration range was limited and these distributions did not include zero rewards and were skewed. We conclude that risk aversion to variability in reward amount is a robust phenomenon for some reward distributions. A systematic evaluation of the effect of various reward distribution parameters on choice behavior should complement functional and mechanistic approaches.

Key words: *Apis mellifera*, coefficient of variation, variance, skew, zero reward, nectar concentration, nectar volume.

Introduction

Foraging alternatives may differ not only in their mean expected value but also in the distribution of reward amounts within each alternative. For example, plant species in a xeric Mediterranean ecosystem (Petanidou and Smets, 1995) differed greatly both in mean nectar amounts and in their variability (Shafir et al., 2003; Weber et al., 2004). Cartar (Cartar, 1991) studied the choice of bumblebees between two co-occurring plants, seabrush (*Plectritis congesta*) and dwarf huckleberry (*Vaccinium caespitosum*), which provided similar mean expected rates of net energy intake but that differed in variability. Risk sensitivity refers to how an animal responds to such variability.

Risk sensitivity can be addressed by both qualitative and quantitative questions. The former focus on the conditions that lead to seeking variability (risk proneness), shying away from variability (risk aversion) or being indifferent (risk insensitivity). According to the energy budget rule, for example, the state of the animal determines whether it will be risk prone or risk averse (Barnard and Brown, 1985; Moore and Simm, 1986; Caraco et al., 1990; Cartar, 1991). The nature of the variable resource also affects choice behavior. Animals are typically risk averse when variability is in a hedonically positive outcome, such as amount of food, and are risk prone when variability is in a hedonically negative outcome, such as delay to receiving food (Kacelnik and Bateson, 1996; Marsh and Kacelnik, 2002).

A quantitative approach focuses on the degree of risk sensitivity (aversion or proneness), or how strong the preference is between two options that differ in variability. For example, honey bee (*Apis mellifera*) workers were more risk averse than drones (Shafir et al., 2005), bumblebees (*Bombus sandersonii*) were more risk averse than wasps (*Vespula vulgaris*) (Real, 1981), and there are cultural differences in humans in levels of risk sensitivity (Weber and Hsee, 2000; Weber et al., 2004). In the present work, we took a quantitative approach, testing degree of risk sensitivity of subjects on a positive energy budget to variability in

reward amount, conditions that are expected to lead to risk aversion.

Our understanding of risk-sensitive choice behavior is best achieved by combining functional, mechanistic and descriptive perspectives (Kacelnik and Bateson, 1996; Shafir, 2000). The functional perspective is concerned with the choice behavior that is predicted to have evolved through natural selection to maximize the animal's fitness, the mechanistic with the process by which choice develops, and the descriptive with what the animal actually chooses. Functional models are needed to explain the effect of state on shifting between risk aversion and proneness, whereas mechanistic models can explain such shifts between positive and negative outcomes (Marsh and Kacelnik, 2002) and differences between experimental conditions in degree of risk sensitivity (Kacelnik and Abreu, 1998; Shapiro, 2000; Shapiro et al., 2001; Shafir et al., 2005).

In their review of risk-sensitivity experiments, Kacelnik and Bateson point out that apparently minor details in experimental design may affect both the presence and direction of risk-sensitive preferences (Kacelnik and Bateson, 1996). In experiments in which variability is in reward amount, the main factors that have been implicated in affecting degree of risk sensitivity are the coefficient of variation (CV=s.d./mean) of the variable alternative (Shafir, 2000), whether or not the variable alternative includes zero (empty) rewards (Perez and Waddington, 1996; Waddington, 1997) and whether variability is in reward volume or concentration (Banschbach and Waddington, 1994; Waddington, 1995; Perez and Waddington, 1996).

A systematic examination of parameters that may affect levels of risk sensitivity is needed. We believe that a good quantitative descriptive model can guide the development of functional and mechanistic models and help in designing better experiments to evaluate them. The goal of this work was to systematically test the effect of CV, the presence of zero rewards and whether variability is in reward volume or concentration on risk sensitivity of honey bees. We tested bees using the proboscis extension response (PER) paradigm (Menzel and Bitterman, 1983), which allows accurate control of stimuli presentation and which Shafir et al. (Shafir et al., 1999; Shafir et al., 2005) have modified for testing risk sensitivity.

Materials and methods

Restraint of subjects

Bees were maintained in standard honey bee hives at the apiary of the B. Triwaks Bee Research Center, Rehovot, Israel, and were free to fly and forage. In order to minimize variability between subjects due to colony genetic and environmental effects, all bees were of the same race, New World Carniolan (Cobey, 1999). All subjects in the odor learning and discrimination experiment were from a single colony, all subjects in the odor preferences experiment and 90% of those in the risk sensitivity experiment were from a second colony, and the rest were from two other colonies.

We harnessed subjects as in Shafir et al. (Shafir et al., 1999; Shafir et al., 2005). Foragers were collected into small glass vials as they returned to the hive. To facilitate harnessing of the bees, each vial was submerged into ice water until the bee stopped moving (~1 min). Subjects were then strapped into a sectioned hollow plastic tube by a 3-mm-wide strip of duct tape that wrapped around the tube and (dorsal) thorax of the bee. The abdomen of the bee was not covered. Subjects were harnessed so that the stand extended to just below the front pair of legs, which were loose over the stand, to ensure that the head of each bee was free to rotate.

After about 30 min, when the bees had recovered from the ice water, we gently squeezed the abdomen of each bee, collected the regurgitated contents of the crop with a micropipette and measured its concentration with a refractometer. Bees that had pollen loads in their curbiculae and that regurgitated fewer than 5 μl were considered pollen collectors, and those without pollen loads and that regurgitated more than 5 μl fluid of at least 10 Brix (to distinguish from water collectors) were considered nectar collectors (Page and Fondrk, 1995).

To avoid starvation, we fed bees 0.8 μl sucrose solution (35% w/w) 1 h after harnessing them. After one more hour, we conducted a motivation test. We gently touched the antennae of each bee with a drop of sucrose solution and only selected bees that extended their proboscides in response to the sugar stimulus; typically, very few bees do not pass this performance criterion.

Apparatus

Odors (conditioned stimulus; CS) were delivered to subjects from 1-ml glass syringes mounted at a training station. We added 3.5 μl of pure odor to a strip of filter paper that was placed inside a syringe. The tip of each syringe was attached by silicone tubing to a valve that was attached to an air pump. Valves were controlled by a computer and opening of a valve caused an odor air stream flow of 13 $cm^3\ s^{-1}$ out of the tip of the syringe and over the subject's antennae. To create an exhaust stream, we connected a 9-cm-diameter tube to a vent and mounted it 13 cm behind the subject. Thus, subjects experienced a constant slow air flow and, when a particular valve was opened, an air stream of the corresponding odor flowed over the antennae of a subject and immediately into the exhaust stream.

Subjects were lined up on a ruler, at 4-cm intervals, with a partition between each subject and its neighbors. After a trial with one subject, the ruler was slid until the next bee was in position, and so forth.

Odor learning and discrimination

Our goal in the learning, discrimination and preference experiments was to choose the odors to be used in the risk sensitivity experiments. The learning and discrimination experiment was conducted at a time when few foragers were collecting pollen, and hence we only tested foragers that returned to the hive without pollen pellets. The experiment

consisted of a training phase, which yielded learning curves for various odors, and a test phase, which involved constructing a discrimination matrix between odors. In the training phase, subjects were conditioned to one of 16 odors. There were six conditioning trials with an intertrial interval of 8 min. A trial began when a bee was placed in the training station. We allowed the subject a few seconds to acclimate and then we presented the odor for 5 s. After 3 s we fed the bee 0.4 μl of a 35% w/w sucrose solution. We noted whether the subject extended its proboscis after the onset of odor delivery but before delivery of the reward. We lightly touched the subject's antennae with the tip of the syringe to induce proboscis extension and the subject was allowed to imbibe the sucrose reward; subjects always ingested the entire droplet. Once a subject learned the association and extended its proboscis after odor presentation, we brought the tip of the syringe directly to the tip of the proboscis.

The test phase was conducted 30 min after the training phase and consisted of two extinction (unrewarded) trials, with an intertrial interval of 8 min. One of the odors tested was always different from the conditioned odor; the other odor tested was either another different odor for some subjects, or the conditioned odor for others. The order of odor presentation was balanced across subjects.

Since the goal of this experiment was to find odors that bees can learn well and discriminate well one from the other, we progressively stopped testing odors that appeared to be inadequate. We eventually concentrated on four odors: 1-octanol (Sigma, St Louis, MO, USA; cat no. 0-4500), benzyl acetate (Aldrich, Milwaukee, WI, USA; cat. no. B01,580-5), eugenol (Merck, Hohenbrunn, Germany; cat no. 8184550100), and geranyl acetate (Aldrich; cat. no. 17,349-5).

Odor preferences

Once we identified four odors that subjects learned and discriminated well, we wanted to choose two of these that are equally preferred by subjects, to be used in the risk-sensitivity choice experiment. The idea was to condition subjects equally to two odors and then test their preference between the odors. Due to differences in sensitivity to stimuli and learning performance between foragers performing different tasks, we conducted the odor preferences and risk-sensitivity experiments only with pollen foragers, which are good learners (Scheiner et al., 2001; Drezner-Levy, 2004; Latshaw and Smith, 2005).

Choice can be tested with the modified PER paradigm (Shafir et al., 1999; Shafir et al., 2005). We attached each of two odor syringes to a base that mounted onto tracks at the training station. Syringes were mounted horizontally so that when we placed a subject in the training station the tips of the syringes were 10 mm from the bee and pointed towards the base of the bee's antennae. Each subject was positioned so that syringes were 30° to the right and the left of its sagittal plane. A line drawn on the base of the station defined the midline between the syringes.

The odor preferences experiment consisted of a training phase and a test phase. The training phase was similar to that of the odor learning experiment above, except that subjects experienced two odors in sequential trials using the sequence ABABABAB. There were eight training trials, four with each odor. We alternated the position (left or right) of each odor every two trials to control for possible side biases. We tested all six combinations of the four odors of interest. For each combination, one odor appeared first for half of the subjects and second for the other half of the subjects.

The test phase was conducted 20 min after the training phase and consisted of presentation of the two odors in an alternating, pulsed schedule. The schedule consisted of 0.8 s of one odor, followed by 0.2 s of no odor, followed by 0.8 s of the other odor, and so forth, until each odor was presented twice. We scored the orientation of the head of each subject with respect to the midline between the two syringes after the last of the four odor pulses, when the computer emitted an audible signal. We scored a choice on every trial, even if the head of a bee showed only a slight deviation from the midline. A video camera mounted above the training station facilitated scoring of choices. The chosen odor was presented for an additional 3.5 s, and the subject was rewarded 1.5 s after the onset of odor. We delivered rewards regardless of whether or not a subject extended its proboscis to the chosen odor. There were four test trials, with an intertrial interval of 9 min.

Risk sensitivity

The goal of the main experiment was to test the effect on risk sensitivity of various parameters that define a variable reward distribution. To reduce the variability between subjects in learning performance, we first conducted a learning phase (similar to the odor learning experiment above) consisting of three trials with eugenol as the conditioned odor. Only subjects that responded to the CS in two or three of the three trials were selected for the risk-sensitivity phase.

The risk-sensitivity phase consisted of choice trials between two odors, as in the test phase of the odor preferences experiment. The two odors were benzyl acetate and geranyl acetate, which were chosen based on the results from the first set of experiments (see Results). For each subject, one odor was associated with a constant reward and the other with a variable reward. The odor assigned to the constant reward was counterbalanced among subjects to control for possible preferences for odors that were not due to associated rewards. We alternated the order in which the two odors were pulsed across trials to control for a possible sequence effect of odor presentation. To control for possible side preferences, we presented each odor on the left (L) or right (R) in the sequence RLLRLRRL…LRRL. The other odor was always presented from the opposite side.

Each subject was tested in one of the experimental conditions in Table 1. In conditions A1–A6, in which variability was in reward volume, reward was administered with a syringe pump (SP200; World Precision Instruments, Sarasota, FL, USA), which allowed the administration of minute amounts accurately. Reward concentration in these

Table 1. *The values of the constant and variable rewards in each condition, and the values of the corresponding mean, variance and coefficient of variation (CV)*

		Variable reward				
Condition	Constant reward	Low (prob.)	High (prob.)	Mean	Variance	CV
A1	0.5	0 (0.8)	2.5 (0.2)	0.5	1	200
A2	1	0 (0.5)	2 (0.5)	1	1	100
A3	0.2	0 (0.8)	1 (0.2)	0.2	0.16	200
A4	0.8	0 (0.8)	4 (0.2)	0.8	2.56	200
A5	0.5	0.1 (0.86)	3 (0.14)	0.5	1	200
A6	1	0.2 (0.61)	2.25 (0.39)	1	1	100
C1	15	5 (0.8)	55 (0.2)	15	400	133
C2	25	5 (0.5)	45 (0.5)	25	400	80
C3	45	5 (0.2)	55 (0.8)	45	400	44

In conditions A1–A6, variability was in volume of reward (μl), and in conditions C1–C3 variability was in concentration of reward (%). The probabilities of low and high variable rewards are in parentheses.

conditions was 35% w/w sucrose solution. For a zero reward, we gently touched the subject's antennae with the tip of a clean syringe and allowed it to touch the empty syringe if it extended its proboscis. In conditions C1–C3, in which variability was in reward concentration, reward was administered with a Gilmont microsyringe, since several syringes (with different concentrations) were needed. Reward amount in these conditions was 0.8 μl.

The sequence of low and high rewards in the variable alternative was predetermined for each condition according to the appropriate probabilities and distributed across the 24 trials in a regular manner. For half of the subjects, the sequence started with high reward, followed by low reward(s), then high again, and so forth, and for the other half of the subjects the sequence started with low reward(s), followed by high reward, then low reward(s) again, and so forth. Every time a subject chose the odor that corresponded to the constant reward it received the appropriate volume (or concentration) of sucrose solution. The first time that a subject chose the variable reward it received the first value in the sequence, and the next time that it chose the variable reward it received the next value in the sequence, and so forth. Thus, every subject within an experimental condition experienced a similar probability as other subjects of low and high values of the variable reward.

We conducted 24 trials with each subject, with an intertrial interval for each subject of 6 min. Up to 12 subjects were tested concurrently every day, randomly assigned to several experimental conditions. We tested all experimental conditions concurrently throughout the duration of the study to avoid possible seasonal biases.

Statistical analyses

In the training phase of the odor learning and discrimination experiment we summed the total number of proboscis extension responses to the CS for each subject over the six conditioning trials as a measure of learning performance. In the test phase, we compared the durations of the proboscis extension response to the odors. The distributions of total number of responses and of proboscis extension durations were not normal, hence we used nonparametric methods. We tested differences between odors using the Wilcoxon and Kruskal-Wallis tests.

In the PER paradigm modified for choice between two odors, subjects exhibit choice in trials in which they extend the proboscis; the orientation of the head when subjects do not respond with proboscis extension is not informative (Shafir et al., 1999; Shafir et al., 2005). Hence, in the odor preference experiment we only considered trials in which subjects responded to the odor stimuli in the test phase, and we calculated a choice proportion between odor pairs for each subject.

In the risk-sensitivity experiment, we performed ANOVA in which the dependent variable was the proportion choice of the constant reward for each subject during the last five trials in which the subject responded to the CS, when choice tended to stabilize. Thirty five of 391 subjects (9%) responded in fewer than five trials and were excluded. The independent variable was experimental condition, and we performed separate analyses for the variability in volume and concentration conditions. We performed multiple comparison post tests using Tukey's method. We also tested the effects of CV and zero rewards (as nominal variables) in a two-way ANOVA for conditions A1, A2, A5 and A6.

Results

Odor learning, discrimination and preference

We wanted to find odors that bees can learn equally well, can discriminate well and yet do not have preferences for. We found four odors for which the spontaneous response during the first trial (prior to conditioning) was relatively low (mean proportion=0.10, N=418 bees) and similar for all odors (Kruskal-Wallis test, χ^2_3=5.1, NS) and which subjects learned equally well (mean proportion response in trial 6=0.72; Kruskal-Wallis test, χ^2_3=0.9, NS).

Subjects discriminated well between the four odors and did

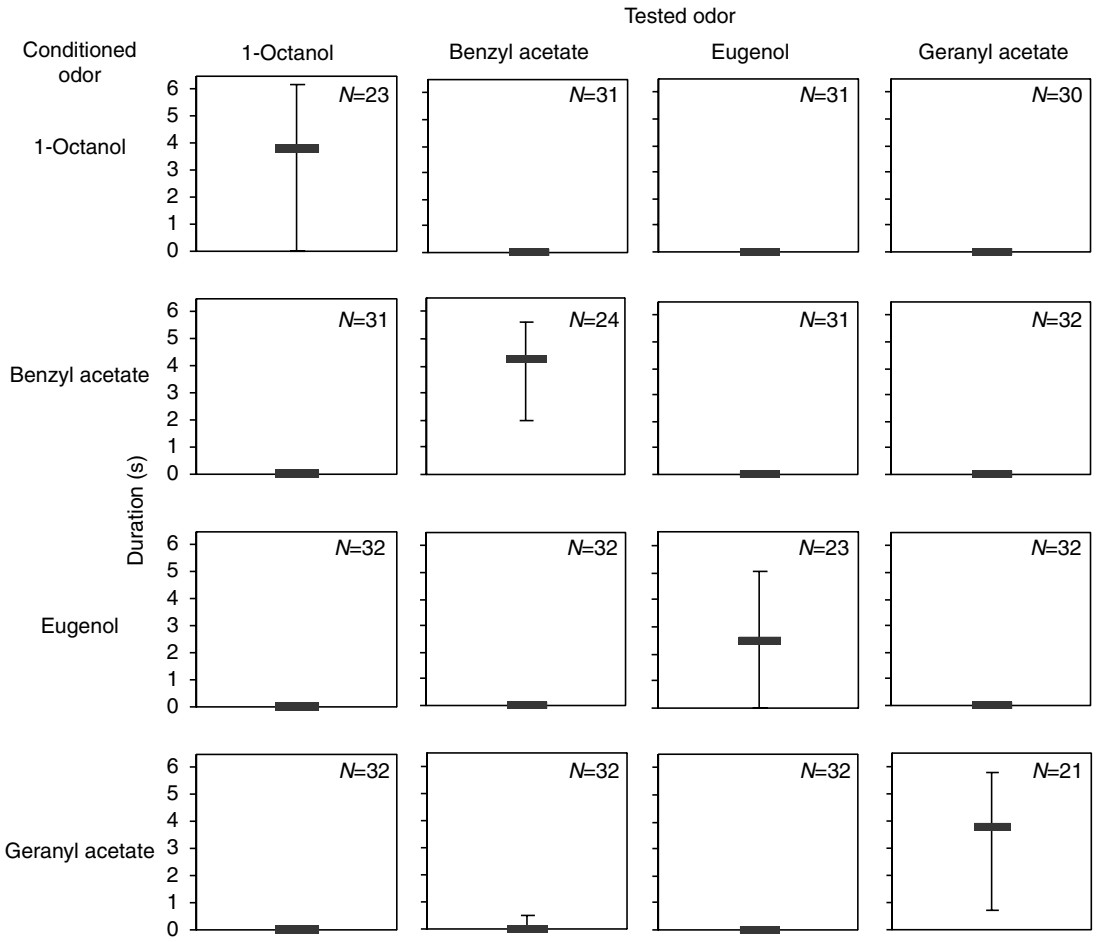

Fig. 1. Generalization matrix between 1-octanol, benzyl acetate, eugenol and geranyl acetate. Each subject was conditioned with one of these odors and then tested with two of them (one of the test odors may have been the same as the conditioned odor). Thick horizontal bars represent median values for the duration of the proboscis extension response during the extinction (unrewarded) test trials. Error bars are the 25 and 75 percentiles (when the tested odor was different from the conditioned odor, these equaled 0 for all but one case). The number of subjects tested in each odor combination is noted.

not generalize between the conditioned odor and the unconditioned odors (Fig. 1). The duration of the proboscis extension response was greater during extinction trials with the conditioned odor than with the unconditioned odors for all four odors (Wilcoxon tests, 1-octanol: χ^2_1=19.5, P<0.001; benzyl acetate: χ^2_1=34.3, P<0.001; eugenol: χ^2_1=22.8, P<0.001; geranyl acetate: χ^2_1=27.4, P<0.001). In fact, subjects hardly responded to the unconditioned odors, with the median duration being zero (no proboscis extension) in all cases, and the 75 percentile also zero in all but one case. The level of response to the three unconditioned odors was similar regardless of the conditioned odor (Kruskal-Wallis tests, 1-octanol: χ^2_2=1.31, NS; benzyl acetate: χ^2_2=3.55, NS; eugenol: χ^2_2=0.53, NS; geranyl acetate: χ^2_2=2.48, NS).

In binary choice tests, subjects generally did not reveal high preference for either of the odors, with mean preferences for the six odor combinations ranging between 0.53 and 0.63. The pair for which preference was closest to 0.5, and for which the 95% confidence interval was smallest, was geranyl acetate and benzyl acetate. These odors were consequently chosen for the risk-sensitivity experiments.

Variability in reward volume

Effect of CV and zeros

In conditions A1, A2, A5 and A6, the variance was 1 μl^2, but the variable option included or did not include zero rewards, and the CV was 200 or 100. A two-way ANOVA testing the effect of zero rewards and CV on the mean proportion choice of constant revealed a significant zeros × CV interaction ($F_{1,157}$=10.8, P=0.001). With zeros, risk aversion was greater for CV=200 (mean proportion choice of constant=0.86) than for CV=100 (0.67) (Fig. 2). Without zeros, subjects were risk insensitive whether the CV=200 (0.49) or CV=100 (0.56).

Effect of variance

Conditions A4, A1 and A3 included zero rewards and the CV was 200, but the variance was 2.56, 1.00 and 0.16 μl^2,

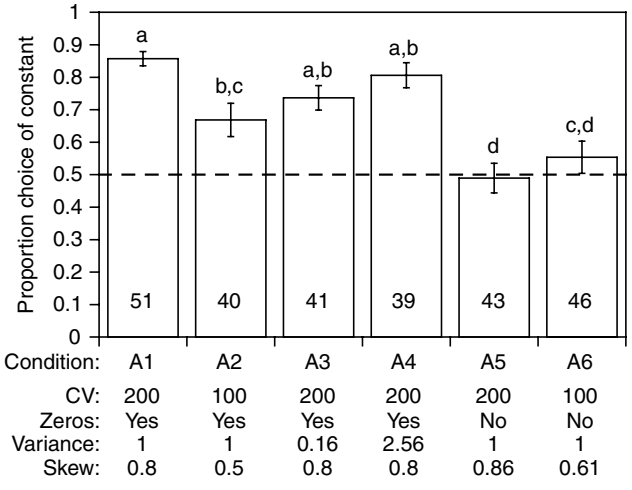

Fig. 2. Mean (± s.e.m.) proportion choice of the constant reward in conditions of the risk-sensitivity experiment in which variability was in reward volume. For each condition, the values of the coefficient of variation (CV) of the variable reward distribution, whether it included zero rewards, the variance (μl^2) and the skew (probability of the low reward) are listed. Distributions are positively skewed for skew >0.5. The broken line represents indifference. Different lowercase letters represent means that are significantly different from each other (P<0.05, Tukey's test). Numbers in the columns are sample sizes.

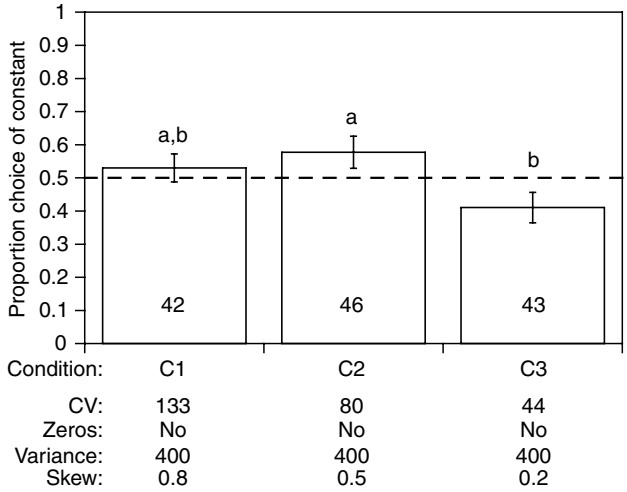

Fig. 3. Mean (± s.e.m.) proportion choice of the constant reward in conditions of the risk-sensitivity experiment in which variability was in reward concentration. For each condition, the values of the coefficient of variation (CV) of the variable reward distribution, whether it included zero rewards, the variance ($\%^2$) and the skew (probability of the low reward) are listed. Distributions are positively skewed for skew >0.5, and negatively skewed for skew <0.5. The broken line represents indifference. Different lowercase letters represent means that are significantly different from each other (P<0.05, Tukey's test). Numbers in the columns are sample sizes.

respectively. Subjects were risk averse in all three conditions, with mean proportion choice ranging between 0.74 and 0.86, and did not differ from each other (Tukey's test) (Fig. 2).

Variability in reward concentration

In conditions C1, C2 and C3, variability was in reward concentration and CV values were 133, 80 and 44, respectively. Mean proportion choice of the constant option ranged between 0.41 and 0.58. In all three conditions, choice proportions did not differ significantly from 0.5 (Wilcoxon signed ranks tests, P>0.05), however proportion choice of constant was significantly greater in condition C2 than in condition C3 (Tukey's test) (Fig. 3).

Discussion

Subjects in all experimental conditions were on a positive energy budget (Shafir et al., 2005) and predicted by the energy budget rule to be risk averse (Stephens and Krebs, 1986). We wanted to test how various parameters that characterize a variable reward distribution affect degree of risk sensitivity. In agreement with meta-analyses of animal (Shafir, 2000) and human (Weber et al., 2004) studies, the CV was a better predictor of risk sensitivity than the variance. In fact, keeping the CV and skewness constant and varying the variance did not affect choice behavior (conditions A1, A3 and A4). Thus, despite the normative appeal in theoretical formulations of using the variance as a measure of variability, it is not a recommended measure for describing choice behavior.

Mean reward (expected value) was not found to affect risk sensitivity when included in the meta-analysis of animal studies (Weber et al., 2004). Our experiments were not specifically designed to test the independent effect of reward mean, but no such effect was apparent. Risk sensitivity did not differ between conditions A2, A3 and A4, although the means ranged between 0.2 and 1 μl. And subjects were less sensitive to risk in conditions A5 and A6 than in the other conditions, although the means were similar: 0.5 and 1 μl, respectively. In summary, neither variance nor mean reward predicts risk sensitivity in isolation. The ratio of the standard deviation and the mean, however, is the CV.

In agreement with previous meta-analyses (Weber et al., 2004), subjects were strongly risk averse when the variable distribution included zero rewards and had a large CV (conditions A1, A3 and A4). Adding to the animal data are experiments with free-flying honey bees where the CV values were 173 and 224, respectively, and the proportions choice of the constant reward were about 0.65 and 0.75, respectively (Shapiro, 2000; Shapiro et al., 2001). This pattern was also found in recent human studies in which subjects experienced the reward probabilities (as is always the case in animal studies), rather than have them described (Hertwig et al., 2004; Weber et al., 2004). It appears that one of the strongest statements regarding risk-sensitive choice behavior is that subjects on a positive energy budget are invariably risk averse to variability in reward amount when the variable reward distribution includes zero rewards and has a large CV. In fact, we are not familiar with any such study in which the CV was >200 and the proportion choice of the constant reward was not

>0.65. It should be noted, however, that in animal studies that meet these criteria, like ours, variability was in reward volume. Although the same pattern was found in experiments with humans, where variability was in monetary payoffs (Hertwig et al., 2004; Weber et al., 2004), it remains to be tested whether it also holds for animal studies in which variability is in reward number (e.g. pellets, seeds) or concentration.

When variability is in reward volume and the variable reward distribution includes zero rewards, the effect of the CV on risk sensitivity is not categorical but graded. Risk sensitivity increases with the CV in such cases when comparing both across (Shafir, 2000) and within (Shafir et al., 2005) studies. Human subjects are similarly affected, especially when having to experience the reward (monetary payoff) distributions (Weber et al., 2004). This finding was supported in the present study by the greater degree of risk aversion in experimental condition A1 than in A2. Thus, the well-supported conclusion of strong risk aversion where the CV is high is probably a special case of a more general and robust pattern of risk sensitivity increasing with the CV.

Meta-analyses of animal studies (Shafir, 2000) and human studies (in which the payoffs and probabilities were described, rather than experienced) (Weber et al., 2004) did not detect a significant effect on choice behavior of whether the variable reward distribution included zero rewards or not. However, when all options are rewarding, the CV tends to be low (Fig. 4), thus making it difficult to ascertain whether low levels of risk sensitivity are due to the lack of zero rewards or to the low CV. We created a distribution (A5) that did not include zero rewards yet had a high CV. Contrary to the predictions of the CV model, subjects were risk indifferent in this experimental condition. However, in order to create such a distribution, we had to increase the skew; the occurrence of the high variable reward was a rare event (P=0.14).

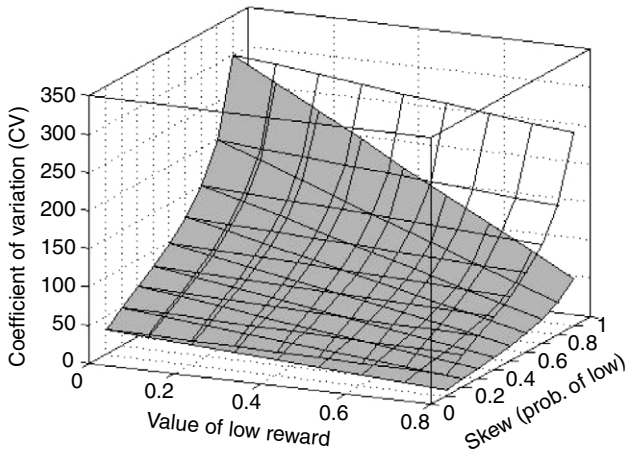

Fig. 4. The effects of the value of the low reward and distribution skew (probability of the low reward) on the coefficient of variation (CV). Calculations are for distributions with probability (P) of low reward value and probability (1–P) of high reward value, when mean reward equals 1 (lower gray surface) or 5 (upper transparent surface).

Distribution skew is known to affect risk sensitivity (Hertwig et al., 2004; Weber et al., 2004). Positively skewed distributions, such as that in condition A5, are predicted to reduce risk aversion, and in fact were associated with reduced risk aversion in the meta-analysis of animal studies (Shafir et al., 2003). Risk aversion was high in conditions A1, A3 and A4 despite relatively high skew. Thus, it appears that the presence of zero rewards and high CV may override the effect of high skew. Also, as for the nonlinear evaluation of probabilities by humans when probabilities and rewards are described (Tversky and Kahneman, 1992), the effect of skew when probabilities and rewards are experienced may be especially important for small probabilities.

Condition A6 had the same CV (=100) as condition A2, but without zero rewards. There was no significant difference in risk sensitivity between the two conditions, supporting the claim that levels of risk sensitivity are affected by the CV and not by the presence or absence of zero rewards. Possibly, lack of zero rewards and a small positive skew in condition A6 may have contributed to a tendency for lower risk sensitivity in that condition relative to A2.

Unlike measures of reward amount, reward concentration is limited to the range 0–100%, and flowers are less variable in nectar concentration than in nectar volume (Shafir et al., 2003). There are conflicting results as to whether pollinators evaluate variability in nectar volume and concentration similarly. In support of this hypothesis, Wunderle and O'Brien concluded that risk aversion in the bananaquits that they studied was affected by the CV of the variable distribution and not by whether variability was in nectar volume or concentration (Wunderle and O'Brien, 1985). In studies with several bee species in which variability was in nectar concentration and the CV=50, subjects were risk indifferent (Banschbach and Waddington, 1994; Waddington, 1995; Perez and Waddington, 1996), similar to choice behavior when variability is in nectar volume and the CV=50. For a variable nectar concentration distribution with greater CV, bumblebees showed greater risk aversion (Waddington, 2001).

However, levels of risk aversion when variability is in nectar concentration are not well described by the CV model. Lack of high levels of risk sensitivity in conditions C1–C3 may be partly due to not having zero rewards in the variable distribution. When variability is in reward concentration, zero rewards consist of water solution (Wunderle and O'Brien, 1985; Shapiro, 2000). We used a 5% concentration for the low reward, which can be detected by honey bees (Frisch, 1950; Afik et al., 2006). We did not find a consistent effect of CV on choice behavior in experimental conditions C1–C3. Whereas risk aversion was greater in condition C2 than in C3, in condition C1 it was lower than expected by its higher CV. Condition C1 was similar to conditions A5 and A6 in having relatively high CV, but no zero rewards, and positive skew. Thus, as in experiments in which variability was in reward volume, the combination of no zero rewards and high positive skew may have reduced risk sensitivity also when variability was in reward concentration.

Also in Shapiro's experiments with free-flying honey bees in which variability was in nectar concentration (Shapiro, 2000), levels of risk aversion were not correlated with the CV. Shapiro was able to simulate with a choice model the behavior of subjects when variability was in either nectar volume or concentration (Shapiro, 2000). However, the model incorporated differently shaped functions for the subjective evaluation of volumes and concentrations. In particular, the curve was linear for concentrations in much of the range, whereas it was concave-down for volumes. Where rewards were evaluated according to such concave-down functions, bees were risk averse regardless of whether variability was in volume (Shafir et al., 2005) or concentration (Waddington, 2001), as explained by Jensen's inequality (Smallwood, 1996). In fact, the effect of the CV on risk sensitivity follows from such concave-down utility functions (Shafir et al., 2003; Weber et al., 2004). Thus, it appears that the CV model is a good predictor of risk sensitivity when evaluation of reward values is described by a concave-down function, which may be more typical of volumes than concentrations.

Levels of risk sensitivity are affected by how subjects perceive the various alternatives, which may lead to intra-specific differences (Shafir et al., 2005) and to differences in sensitivity to variability in volume and concentration (Shapiro, 2000; Shapiro et al., 2001; Waddington, 2001). Nevertheless, some generalizations can be made. We conclude that risk sensitivity to variability in reward amount is more robust than has been previously appreciated (Kacelnik and Bateson, 1996; Marsh and Kacelnik, 2002), at least for some reward distributions. A better understanding of the characteristics of such distributions can be helpful in designing and interpreting risk sensitivity experiments. For example, risk indifference exhibited by starlings (Marsh and Kacelnik, 2002) was probably due to the variable distribution not including zero rewards and having a CV of 50, and not to variability being in reward amount rather than delay. Similarly, risk indifference exhibited by carpenter bees under both negative and positive energy budgets was probably due to the variable distribution not including zero rewards and having a CV of 50 (Perez and Waddington, 1996); the energy budget rule and differences in risk sensitivity between solitary and social foragers should be tested with variable reward distributions to which subjects are expected to be more sensitive. For variability in reward amount, highest levels of risk aversion are found when variability is in volume and the distribution includes zero rewards and has a high CV. The variance itself does not affect risk sensitivity.

Because the CV is a ratio, it is dimensionless and does not depend on scale. That is, if every reward is increased by the same proportion, the CV remains constant and risk sensitivity is predicted to remain the same. Such scale invariance is a common property of many vertebrate conditioning phenomena (Gallistel and Gibbon, 2000). Thus, our findings provide further support for similarity between invertebrate and vertebrate learning (Bitterman, 1996).

Changing the value of a particular parameter describing a variable distribution affects the values of other parameters. For example, greatest CV values are achieved with variable distributions that include zero rewards and are highly positively skewed (Fig. 4). Maintaining constant skew and increasing the value of the low reward decreases the CV; the rate of reduction is faster when mean reward is smaller, increasingly so the greater the skew. Similarly, decreasing skew while maintaining the value of the low reward constant decreases the CV; however, the rate of reduction is faster when mean reward is larger, increasingly so the greater the value of the low reward. Thus, a more fine-grain analysis of the relative independent contribution of various distribution parameters to risky choice would require a multivariate analysis of choice experiments covering the full parameter values space. In particular, such analysis could resolve the relative contribution to risk sensitivity of distribution skew and CV and whether the effect of increasing the value of the low reward is continuous or whether there is a special effect of increasing the value of the low reward to above zero.

This research was supported by the Israel Science Foundation (ISF#513/01). We thank Hillary Voet for statistical advice, and two anonymous referees for valuable comments.

References

Afik, O., Dag, A., Kerem, Z. and Shafir, S. (2006). Analyses of avocado (*Persea americana*) nectar properties and their perception by honey bees (*Apis mellifera*). *J. Chem. Ecol.* **32**, 1949-1963.

Banschbach, V. S. and Waddington, K. D. (1994). Risk-sensitive foraging in honey bees: no consensus among individuals and no effect of colony honey stores. *Anim. Behav.* **47**, 933-941.

Barnard, C. J. and Brown, C. A. J. (1985). Risk-sensitive foraging in common shrews (*Sorex araneus* L.). *Behav. Ecol. Sociobiol.* **16**, 161-164.

Bitterman, M. E. (1996). Comparative analysis of learning in honeybees. *Anim. Learn. Behav.* **24**, 123-141.

Caraco, T., Blanckenhorn, W. U., Gregory, G. M., Newman, J. A., Recer, G. M. and Zwicker, S. M. (1990). Risk-sensitivity: ambient temperature affects foraging choice. *Anim. Behav.* **39**, 338-345.

Cartar, R. V. (1991). A test of risk-sensitive foraging in wild bumble bees. *Ecology* **72**, 888-895.

Cobey, S. (1999). The New World Carniolan closed population breeding project. In *Proceedings of the 36th Apimondia Congress, Vancouver, Canada*, pp. 26-27. Bucharest: Apimondia Publishing House.

Drezner-Levy, T. (2004). Evaluation of reward distribution by honey bee (*Apis mellifera*) foragers collecting nectar, pollen and water. Thesis, Faculty of Agricultural, Food and Environmental Quality Sciences. Rehovot: Hebrew University of Jerusalem.

Frisch, K. V. (1950). Bees – their vision, chemical senses, and language. Ithaca, NY: Cornell University Press.

Gallistel, C. R. and Gibbon, J. (2000). Time, rate, and conditioning. *Psychol. Rev.* **107**, 289-344.

Hertwig, R., Barron, G., Weber, E. U. and Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.* **15**, 534-539.

Kacelnik, A. and Abreu, F. B. E. (1998). Risky choice and Weber's law. *J. Theor. Biol.* **194**, 289-298.

Kacelnik, A. and Bateson, M. (1996). Risky theories – the effects of variance on foraging decisions. *Am. Zool.* **36**, 402-434.

Latshaw, J. S. and Smith, B. H. (2005). Heritable variation in learning performance affects foraging preferences in the honey bee (*Apis mellifera*). *Behav. Ecol. Sociobiol.* **58**, 200-207.

Marsh, B. and Kacelnik, A. (2002). Framing effects and risky decisions in starlings. *Proc. Natl. Acad. Sci. USA* **99**, 3352-3355.

Menzel, R. and Bitterman, M. E. (1983). Learning by honeybees in an

unnatural situation. In *Neuroethology and Behavioral Physiology* (ed. F. Huber and H. Markl), pp. 206-215. Berlin, Heidelberg: Springer-Verlag.

Moore, F. R. and Simm, P. A. (1986). Risk-sensitive foraging by a migratory bird (*Dendroica coronata*). *Experientia* **42**, 1054-1056.

Page, R. E. and Fondrk, M. K. (1995). The effects of colony level selection on the social-organization of honey-bee (*Apis-mellifera* L) colonies – colony level components of pollen hoarding. *Behav. Ecol. Sociobiol.* **36**, 135-144.

Perez, S. M. and Waddington, K. D. (1996). Carpenter bee (*Xylocopa micans*) risk indifference and a review of nectarivore risk-sensitivity studies. *Am. Zool.* **36**, 435-446.

Petanidou, T. and Smets, E. (1995). The potential of marginal lands for bees and apiculture – nectar secretion in Mediterranean shrublands. *Apidologie* **26**, 39-52.

Real, L. A. (1981). Uncertainty and pollinator-plant interactions: the foraging behavior of bees and wasps on artificial flowers. *Ecology* **62**, 20-26.

Scheiner, R., Page, R. E. and Erber, J. (2001). The effects of genotype, foraging role, and sucrose responsiveness on the tactile learning performance of honey bees (*Apis mellifera* L.). *Neurobiol. Learn. Mem.* **76**, 138-150.

Shafir, S. (2000). Risk-sensitive foraging: the effect of relative variability. *Oikos* **88**, 663-669.

Shafir, S., Wiegmann, D. D., Smith, B. H. and Real, L. A. (1999). Risk-sensitive foraging: choice behaviour of honeybees in response to variability in volume of reward. *Anim. Behav.* **57**, 1055-1061.

Shafir, S., Bechar, A. and Weber, E. U. (2003). Cognition-mediated coevolution – context-dependent evaluations and sensitivity of pollinators to variability in nectar rewards. *Plant Syst. Evol.* **238**, 195-209.

Shafir, S., Menda, G. and Smith, B. H. (2005). Caste-specific differences in risk sensitivity in honeybees, *Apis mellifera*. *Anim. Behav.* **69**, 859-868.

Shapiro, M. S. (2000). Quantitative analysis of risk sensitivity in honeybees (*Apis mellifera*) with variability in concentration and amount of reward. *J. Exp. Psychol. Anim. Behav. Process.* **26**, 196-205.

Shapiro, M. S., Couvillon, P. A. and Bitterman, M. E. (2001). Quantitative tests of an associative theory of risk-sensitivity in honeybees. *J. Exp. Biol.* **204**, 565-573.

Smallwood, P. D. (1996). An introduction to risk sensitivity: the use of Jensen's inequality to clarify evolutionary arguments of adaptation and constraint. *Am. Zool.* **36**, 392-401.

Stephens, D. W. and Krebs, J. R. (1986). *Foraging Theory*. Princeton: Princeton University Press.

Tversky, A. and Kahneman, D. (1992). Advances in prospect-theory – cumulative representation of uncertainty. *J. Risk Uncertain.* **5**, 297-323.

Waddington, K. D. (1995). Bumblebees do not respond to variance in nectar concentration. *Ethology* **101**, 33-38.

Waddington, K. D. (1997). Foraging behavior of nectarivores and pollen collectors. *Acta Hortic.* **437**, 175-191.

Waddington, K. D. (2001). Subjective evaluation and choice behavior by nectar- and pollen-collecting bees. In *Cognitive Ecology of Pollination: Animal Behavior and Floral Evolution* (ed. L. Chittka and J. D. Thompson), pp. 41-60. Cambridge: Cambridge University Press.

Weber, E. U. and Hsee, C. K. (2000). Culture and individual judgment and decision making. *Appl. Psychol. Int. Rev.* **49**, 32-61.

Weber, E. U., Shafir, S. and Blais, A. R. (2004). Predicting risk sensitivity in humans and lower animals: risk as variance or coefficient of variation. *Psychol. Rev.* **111**, 430-445.

Wunderle, J. M. and O'Brien, T. G. (1985). Risk aversion in hand-reared bananaquits. *Behav. Ecol. Sociobiol.* **17**, 371-380.