Does risk-sensitivity transfer across movements?

Megan K. O'Brien and Alaa A. Ahmed

J Neurophysiol 109:1866-1875, 2013. First published 16 January 2013;
doi: 10.1152/jn.00826.2012

J Neurophysiol 109: 1866–1875, 2013.
First published January 16, 2013; doi:10.1152/jn.00826.2012.

Does risk-sensitivity transfer across movements?

Megan K. O'Brien[1] and Alaa A. Ahmed[1,2]
[1]*Department of Mechanical Engineering, University of Colorado Boulder, Boulder, Colorado; and* [2]*Department of Integrative Physiology, University of Colorado Boulder, Boulder, Colorado*

Submitted 17 September 2012; accepted in final form 13 January 2013

O'Brien MK, Ahmed AA. Does risk-sensitivity transfer across movements? *J Neurophysiol* 109: 1866–1875, 2013. First published January 16, 2013; doi:10.1152/jn.00826.2012.—An intriguing finding in motor control studies is the marked effect of risk on movement decision making. However, there are inconsistent reports of risk-sensitivity across different movements and tasks, with both risk-seeking and risk-averse behavior observed. This raises the question of whether risk-sensitivity in movement decision making is context dependent and specific to the movement or task being performed. We investigated whether risk-sensitivity transfers between dissimilar movements within a single task. Healthy young adults made arm-reaching movements or whole-body leaning movements to move a cursor as close to the edge of a virtual cliff as possible without moving beyond the edge. They received points on the basis of the cursor's final proximity to the cliff edge. Risk was manipulated by increasing the point penalty associated with the cliff region and/or adding Gaussian noise to the cursor. We compared subjects' movement endpoints with endpoints predicted by a subject-specific, risk-neutral model of movement planning. Subjects demonstrated risk-seeking behavior in both movements that was consistent across risk environments, moving closer to the cliff than the model predicted. However, subjects were significantly more risk-seeking in whole-body movements. Our results present the first evidence of risk-sensitivity in whole-body movements. They also demonstrate that the direction of risk-sensitivity (i.e., risk-seeking or risk-averse) is similar between arm-reaching and whole-body movements, although degree of risk-sensitivity did not transfer from one movement to another. This finding has important implications for the ability of quantitative descriptions of decision making to generalize across movements and, ultimately, decision-making contexts.

sensorimotor control; decision making; arm movement; whole-body movement; biomechanics

EACH ONE OF OUR MOVEMENTS represents a decision made under risk. The potentially injurious consequences of a poor decision compel us to investigate the mechanisms underlying the decision process. This process is partly determined by our sensitivity to risk, i.e., whether one is risk-seeking or risk-averse. However, it remains unclear whether an individual maintains the same risk-sensitivity across movements. Will a risk-averse skier also be risk-averse when playing golf? Can we recover an individual's risk-sensitivity in one task by simply observing his/her behavior in another?

A statistical framework is often used to examine movement decision making in a variety of tasks that involve risk (Faisal et al. 2008; Trommershäuser et al. 2003, 2008; van Beers et al. 2002). Models of optimal movement planning determine how an individual should behave to maximize expected reward, accounting for extrinsic costs and sensorimotor variability. In past studies, these models have correctly predicted the behavior of human subjects in goal-directed pointing movements with a symmetric expected gain landscape. This suggests that subjects performed optimally in these movements, accurately internalizing their own sensorimotor variability and additional task-related costs (Trommershäuser et al. 2003). The high degree of optimality also indicates that subjects adopted a risk-neutral attitude during this motor task (Braun et al. 2011). In other words, they were not sensitive to the environmental risk, defined here as the variance over potential outcomes. Contrastingly, risk-sensitive behavior may emerge if an individual considers both risk and return when deciding how to act, such as proposed by Markowitz (1952) in a financial setting. Individuals may also manifest risk-sensitivity if they are unable to appropriately evaluate the reward structure of the task (distorted utility weighting) or have a skewed weighting of their sensorimotor variability (distorted probability weighting). Other movement studies have found evidence of risk-seeking behavior (Wu et al. 2009) or risk-averse behavior (Nagengast et al. 2010). However, this previous work has examined different types of movement (pointing, arm-reaching) in a variety of experimental paradigms, has not always provided feedback immediately after the movement, and has involved comparisons between subjects. In the present work, we investigate movement decisions in a single paradigm, providing the same form of feedback but with two different movements to determine whether risk-sensitivity transfers from one type of movement to another.

The two movements we chose to compare were arm-reaching and whole-body leaning movements. Although many studies examine arm-reaching, risk is arguably more relevant to whole-body movements. Furthermore, goal-directed whole-body movements are less familiar than arm-reaching. If risk-sensitivity transferred between movements, this would be a strong demonstration of generalization. If risk-sensitivity did not transfer, this would establish its dependence on movement context.

We designed an experiment to investigate risk-sensitivity in arm-reaching and whole-body movements. We manipulated risk in the form of point penalties and/or sensorimotor variability, thereby allowing us to examine the consistency of movement decisions across various risk environments and to determine whether underlying risk-sensitivity arises from distortions in utility or probability. We hypothesized that subjects would demonstrate risk-sensitive behavior in both movement tasks and that direction and degree of risk-sensitivity would transfer between the tasks. That is, if a subject was risk-seeking in arm-reaching, we expected them to be equally risk-seeking in the whole-body task. Our findings may advance an under-

Address for reprint requests and other correspondence: A. Ahmed, Univ. of Colorado Boulder, Dept. of Integrative Physiology, 354 UCB, Boulder, CO 80309-0354 (e-mail: alaa.ahmed@colorado.edu).

standing of whether risk-sensitivity in one context can predict movement behavior in other contexts.

MATERIALS AND METHODS

We examined a paradigm in which subjects guided a cursor toward the edge of a virtual cliff using either arm-reaching movements or whole-body leaning movements. Risk was manipulated experimentally in a series of conditions by increasing point penalties and/or task-relevant cursor variability, and subject performance was compared with estimates of a risk-neutral movement planner. Such comparisons allow us to quantify an individual's risk-sensitivity on the basis of their own sensorimotor variability and experimentally imposed risk. Manipulating risk in this task enables us to examine the consistency of the direction and degree of risk-sensitivity. The direction of risk-sensitivity refers to the classification of behavior as risk-neutral, risk-seeking, or risk-averse. The degree of risk-sensitivity refers to the strength, or magnitude, of this preference.

Subjects. Twenty right-handed, healthy subjects (12 men, 8 women; mean age 23.9 ± 2.6 yr) participated in both an arm-reaching and a whole-body movement task. All subjects provided informed consent, and the experimental protocol was approved by the Institutional Review Board of the University of Colorado Boulder.

Experimental protocol. In the arm-reaching experiment (ARM task), subjects used their dominant arm to grasp the handle of a robotic manipulandum (Interactive Motion Technologies Shoulder-Elbow Robot 2) and drive a cursor to the edge of a virtual cliff. In the whole-body experiment (WB task), subjects stood on a force plate (AMTI LG-4-6-1) and used their center of pressure to move the cursor to the cliff edge. A monitor mounted in front of the subject displayed a cursor, a starting position, and penalty region (cliff) set at two-thirds of the subject's maximum movement distance, as shown in Fig. 1. To determine the maximum movement distance, we asked subjects to reach or lean as far forward as they could, and we used the maximum value of three attempts. For the ARM task, subjects were seated and secured with a four-point seatbelt, which prevented rotation of the shoulders or trunk during reaching. For the WB task, subjects wore socks and stood with their feet shoulder-width apart; they also kept their heels in contact with the force plate and their arms crossed in front of their chest. Visual feedback of the arm or body was not intentionally obscured in any way.

We created small visual distinctions between the ARM and WB experiments to encourage subjects to formulate movement strategies independently for each task. One such distinction was the shape of the cursor: in the ARM task, the cursor was a circle of radius 0.25-cm, whereas in the WB task, the cursor was a 0.25 × 2-cm rectangle. Since maximum center-of-pressure movements are inherently smaller than maximum arm-reaching distances, cursor feedback was scaled 2:1 in the WB task. Furthermore, the cliff was not located at the same place on the screen in the ARM and WB experiments to ensure that subjects did not simply aim for the same point on the screen during each movement. The order of movement tasks was varied across subjects. Eleven of the 20 subjects performed the ARM task first, followed by the WB task.

Subjects were instructed to make a swift out-and-back movement to rapidly move the cursor as close to the edge of the cliff as possible without going into the cliff region and rapidly return to their starting position. They received a point score for each trial based on the cursor's maximum excursion to the cliff edge. On the safe side of the cliff, points were awarded as a linear function of movement distance (G_{safe}), and the maximum possible score of 100 points was associated with moving the center of the cursor perfectly to the edge of the cliff. A different score (G_{cliff}, either 0 points or −500 points) was given if the cursor crossed the cliff edge at any time. Before each trial, the subject had to center the cursor in the starting position, and an auditory tone signaled the beginning of a trial. During a trial, the cursor was constrained to move toward and away from the cliff, which corresponded to the subject's anteroposterior direction. In both tasks, cursor movement was one-dimensional and unaffected by side-to-side movement of the arm or center of pressure. Subjects were given 800 ms to complete a trial, and the movement endpoint was taken as the maximum distance moved toward the cliff from the starting position. We imposed this short trial length to discourage subjects from "hovering" near the cliff region and making small adjustments to increase their scores. Rather, subjects had to make a quick movement decision and return to the starting position.

Subjects completed this paradigm under various risky environments. Risk was manipulated by increasing sensorimotor variability and/or point penalties. We tested four risk conditions for each movement task, with 120 trials performed for each condition: *1*) NULL, where there was no point reward or penalty ($G_{\text{cliff}} = 0$ points) for entering the cliff region; *2*) NOISE, where Gaussian

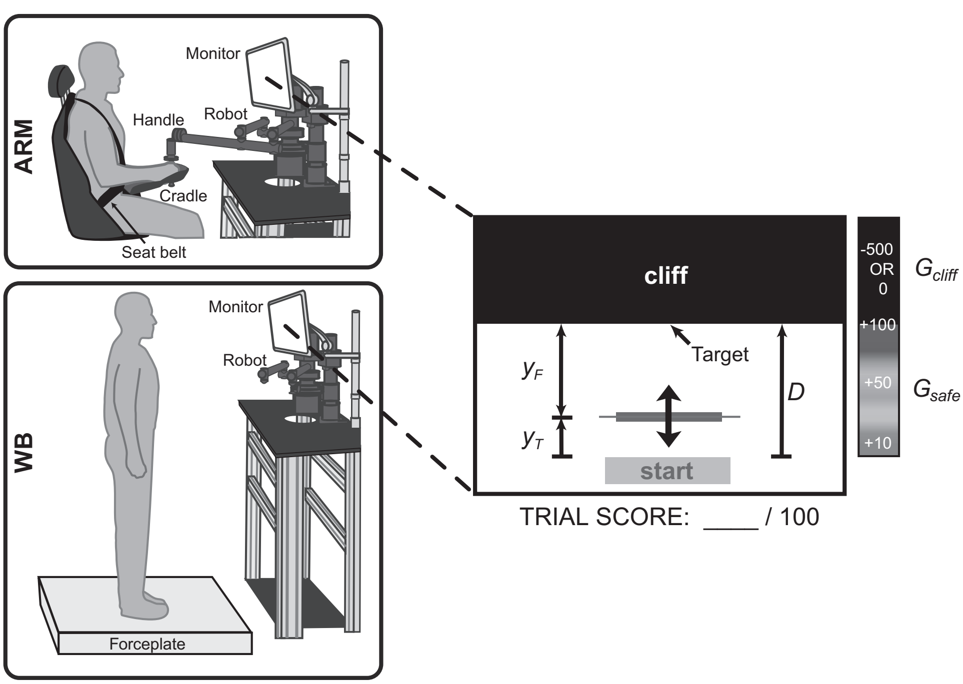

Fig. 1. Experimental setup. *Left*: schematic of arm-reaching (ARM) and whole-body (WB) movement tasks. *Right*: visual feedback of cliff paradigm includes starting position, rectangular cursor, and cliff. Arrow indicates direction of cursor movement. Cursor endpoint determines trial score, either G_{safe} or G_{cliff}. Endpoints are denoted as y_T when referring to a distance traveled toward the cliff and as y_F when referring to a distance from the cliff. The cliff distance D is set at two-thirds of the subject's maximum movement distance.

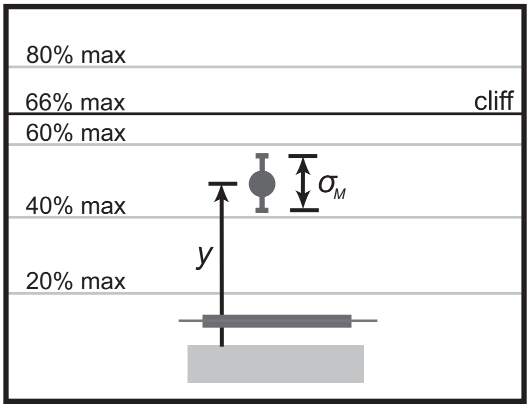

Fig. 2. Variability testing. Visual feedback and variability, σ_M, associated with a mean movement endpoint, y, during variability testing (PRE and POST) is shown. Subjects moved a cursor to 5 discrete targets placed at percentages of their maximum movement distance (max), indicated by the thin horizontal gray and black lines.

noise (with variance $\sigma_N = 0.3$ cm) was added to the cursor position in the direction of movement; *3*) CLIFF, where a penalty ($G_{cliff} = -500$ points) was incurred for moving into the cliff region; and *4*) CLIFF+NOISE, where both the large point penalty and Gaussian noise were included. Adding noise to the cursor and applying a large penalty to the cliff region corresponds to increasing risk by increasing the variability and penalty associated with the movement task, respectively. Thus the tasks are categorized as low penalty and low variability for NULL, low penalty and high variability for NOISE, high penalty and low variability for CLIFF, and high penalty and high variability for CLIFF+NOISE. The cursor noise was perceivable to the eye and was applied for the duration of the out-and-back movement. Subjects also received verbal instructions before each condition describing the value of the cliff penalty. Twelve subjects performed the conditions in the above order (NULL, NOISE, CLIFF, CLIFF+NOISE). To determine whether the order of conditions affected risk-sensitivity, eight subjects performed these conditions in a randomized order.

Variability testing. We determined each subject's sensorimotor variability as a function of movement distance in a separate experiment. Here, we measured each subject's endpoint variability at 5 discrete distances: 20, 40, 60, and 80% of their maximum movement, as well as the distance to the cliff edge (66%), with 40 trials performed for each prescribed distance (Fig. 2). Subjects were again given 800 ms to complete each trial. They completed this variability task twice: once before the four conditions (PRE), and again after completing the four conditions (POST). However, we did not introduce the 66% distance or the POST testing until after the first four subjects had participated in the experiment.

With a known movement distance y (mean endpoint when aiming for each discrete line), target width σ_M (variability of movement endpoints), and the corresponding movement time t, we can use an adapted relation of Fitts' law (Faisal and Wolpert 2009) to estimate two subject-specific parameters, c and d:

$$\sigma_M(y) = y2^{\left(1-\frac{t(y)-c}{d}\right)}. \tag{1}$$

With this approach, we are able to estimate an individual's movement variability as a function of distance, where the movement time is fit as a linear function of movement distance, $t(y) = ay + b$, from the discrete line endpoint data. The only parameters needed as inputs to the statistical decision theory (SDT) model were the four parameters a, b, c, and d, which we averaged from the PRE and POST tasks.

Data acquisition. In the ARM task, optical encoders sampled the position of the robot handle at 200 Hz. In the WB task, the forceplate recorded three-dimensional forces (F_x, F_y, F_z) and moments (M_x, M_y, M_z) about its center at 200 Hz. Center of pressure (COP) was calculated as $[COP_x\ COP_y] = [M_x\ M_y]/F_z$, where x and y refer to mediolateral and anteroposterior axes, respectively.

Risk-sensitivity. We quantified risk-sensitivity by comparing subjects' actual movement endpoints with endpoints predicted by a risk-neutral model of movement planning. We extended a model based on principles of SDT (Trommershäuser et al. 2003) to calculate the risk-neutral movement endpoint for each subject in all four conditions. The model-predicted endpoint is dependent on a subject's sensorimotor variability (determined from our variability test) and the reward/penalty structure of the task.

In this model, a subject's expected gain function $\Gamma(S)$ for a chosen movement strategy S is a product of the probability of hitting a region R_i given that endpoint and the gain G_i associated with the region. In our experiment, the two regions of interest are the safe region, R_{safe} (linear gain G_{safe} ranging from 0 to 100 points), and the cliff region, R_{cliff} (gain G_{cliff} of 0 or -500 points). The expected gain as a function of movement distance is

$$\Gamma(y) = \begin{cases} G_{safe}P(y'|y) & \text{if } y' \le y_{cliff} \\ G_{cliff}P(y'|y) & \text{if } y' > y_{cliff} \end{cases}. \tag{2}$$

We compute the probabilities $P(y'|y)$ by assuming that the actual movement endpoints, y', are distributed around a planned endpoint, y, according to a Gaussian distribution and integrating over the entire landscape ($R_{safe} + R_{cliff}$).

$$p(y'|y) = \frac{1}{2\pi\sigma^2}\exp\left\{\left[-(y'-y)^2\right]/2\sigma^2\right\} \tag{3a}$$

$$P(y'|y) = \int_{R_{safe}+R_{cliff}} p(y'|y)dy' \tag{3b}$$

The total variance σ is a combination of the noise added to the cursor and the subject's sensorimotor variability:

$$\sigma^2 = \sigma_N^2 + [\sigma_M(y)]^2. \tag{4}$$

With a subject's sensorimotor variability as a function of movement distance, this model computes the risk-neutral movement endpoint, or the endpoint that would maximize the expected number of points for a given condition in our cliff paradigm. This endpoint is further from the cliff edge under conditions of increased risk, introduced with either increased penalty and/or variability, as illustrated in Fig. 3.

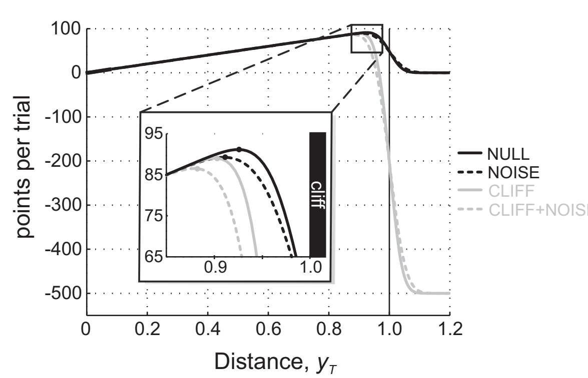

Fig. 3. Risk-neutral movement planner. Sample model-predicted movement endpoints and expected gain landscape during the 4 conditions (NULL, NOISE, CLIFF, CLIFF+NOISE) are shown. The risk-neutral movement endpoint, y_T^{MEG}, maximizes the expected number of points per trial and is denoted by a filled circle on each condition's gain landscape. This endpoint recedes farther from the cliff edge in conditions of increased risk: greater penalty (500-point cliff penalty) and/or variability ($\sigma_N = 0.3$ cm). See text for description of the 4 conditions.

Risk-sensitivity is calculated from the ratio between a subject's mean endpoint, y_F, and model-predicted endpoint, y_F^{MEG}, for each risk condition:

$$\text{risk-sensitivity (\%)} = 100\left(\frac{y_F}{y_F^{MEG}} - 1\right). \tag{5}$$

Thus a risk-sensitivity of 0% indicates perfect agreement between the model prediction and the subject behavior (risk-neutral). A positive risk-sensitivity value indicates that a subject moved farther than the model predicted (risk-seeking), and a negative value indicates that a subject did not move as far as the model predicted (risk-averse).

In this experiment, we can not only classify an individual as risk-sensitive or risk-neutral by movement decisions in any given risk condition, but we can also compare movement across the different conditions to examine the consistency of risk-sensitivity. Comparing movement decisions across conditions may also allow us to determine whether risk-sensitive behavior could be explained by subject-specific distortions in the utility or probability weightings relevant to the virtual cliff paradigm.

We performed 3 separate analyses on all 20 subjects to investigate risk-sensitivity in each task and the transfer of risk-sensitivity between tasks. We repeated these analyses separately for the eight subjects who completed the conditions in a random order. First, we quantified the direction and degree of risk-sensitivity at the group level for each condition and movement task. From this analysis we could determine whether subjects exhibited consistent direction and degree of risk-sensitivity in each movement task. If risk-sensitivity was consistent across conditions, we could draw meaningful conclusions about subjects' overall risk-sensitivity for each movement. Importantly, this analysis would not detect transfer of risk attitudes if individual subjects were idiosyncratic in their risk-sensitivity, but maintained this risk-sensitivity across tasks.

To allow for different risk-sensitivities across subjects, yet still examine consistency across tasks, we turned to a second analysis. Using the same risk-sensitivity data, we performed a subject-level analysis to determine whether the direction of risk-sensitivity was consistent between movements. If a subject was risk-seeking in the arm-reaching movement, was he similarly risk-seeking in the whole-body movement? This involved paired comparisons between risk-sensitivity measures in each movement for each condition and was verified with a regression analysis at the group and subject levels.

Third, to determine possible mechanisms underlying any observed risk-sensitivity, we fit subject-specific risk-sensitivity parameters and compared these parameters between movements. We adjusted our model to incorporate principles of cumulative prospect theory (CPT) introduced by Tversky and Kahneman (1992). In CPT, risk-sensitivity can be explained by a distortion in either the *1*) utility function or *2*) probability weighting function. This leads to an adjusted expected gain function $\Gamma'(y)$ that includes a utility and probability weighting function $w(P)$, similar to that described by Wu et al. (2009):

$$\Gamma'(y) = \begin{cases} G_i^{\alpha} w(P), & \text{if } G_i \geq 0 \\ -(-G_i)^{\beta} w(P), & \text{if } G_i < 0 \end{cases} \tag{6a}$$

$$w(P) = \exp\left\{-\left[-\log P(R_i|y)\right]^{\gamma}\right\} \tag{6b}$$

In our experiment, distorted utility means inappropriately valuing the point rewards and penalties represented by coefficients α and β, respectively. An appropriate utility weighting then corresponds to $\alpha = \beta = 1.0$. Given that many subjects demonstrated a risk-seeking behavior in the movement tasks, we would expect most subjects to overvalue rewards ($\alpha > 1.0$) and undervalue penalties ($\beta < 1.0$), which would likely result in moving closer to the cliff. An appropriate probability weighting corresponds to $\gamma = 1.0$, whereas $\gamma > 1.0$ would represent the overweighting of large probabilities and the underweighting of small probabilities, thus manifesting risk-seeking behavior. We fit values of α, β, and γ for each of our subjects to determine whether distorted weighting of utility or probability could explain any observed risk-sensitivity. Specifically, we used the fminsearch function in MATLAB to find the weighting values that minimized the mean squared error between a subject's actual endpoints and the model predictions in all conditions. This function calculates the local minimum of an unconstrained multivariate function around an initial estimate ($\alpha = 1.0$, $\beta = 1.0$, $\gamma = 1.0$).

Statistics. We performed a three-way repeated-measures analysis of variance (ANOVA) to determine whether there were effects of risk condition (including penalty and variability) or task on movement endpoints. We used independent *t*-tests to compare risk-sensitivity measures to 0% (direction of risk-sensitivity), and we used paired *t*-tests to compare risk-sensitivity measures between the ARM and WB tasks (degree of risk-sensitivity). In the latter case, we adjusted for multiple comparisons across the four risk conditions. For the group-level least-squares regression analysis, we verified statistical evidence for a linear relationship between ARM and WB risk-sensitivity with an *F*-test. We used an independent *t*-test to compare individual subjects' regression slopes to unity. We used paired *t*-tests for subject-specific comparisons between tasks, conditions, and CPT parameters. To adjust for multiple comparisons across risk conditions, we used the Bonferroni correction, with the significance level set to 1.25%. For all other statistical tests, the significance level was set to 5%.

RESULTS

Overview. We found that most subjects' endpoints followed the general trends predicted by a subject-specific model and risk-neutral movement planner. In both ARM and WB tasks, subjects avoided the cliff more with increasing point penalties and increasing noise. Overall, the direction of risk-sensitivity was consistent between the ARM and WB tasks, but the degree of risk-sensitivity did not transfer between these two movement types. These results hold when the order of conditions is randomized.

Movement trends. The mean distance from the starting position to the cliff edge was 15.4 ± 3.2 cm for the ARM task and 5.9 ± 1.4 cm for the WB task. We examined the last 100 trials of each condition. Movement endpoints for a representative subject (S8) performing the ARM and WB cliff tasks are shown in Fig. 4*A*, with all endpoints normalized to that subject's cliff distance. We denote a movement endpoint as y_T when referring to a distance traveled toward the cliff and as y_F when referring to a distance from the cliff ($y_F > 0$ corresponds to movements on the safe side of the cliff). The distribution of endpoints was approximately Gaussian for all subjects. On average, during these last 100 trials, subjects moved past the cliff edge the following number of times: in ARM, 8.5 ± 5.9 (NULL), 1.8 ± 2.5 (NOISE), 2.9 ± 2.7 (CLIFF), and 0.8 ± 1.2 (CLIFF+NOISE); in WB, 17.2 ± 7.6 (NULL), 8.0 ± 5.8 (NOISE), 6.1 ± 4.8 (CLIFF), and 2.6 ± 2.4 (CLIFF+NOISE).

Increasing penalty and variability significantly affected movement endpoints in both the ARM and WB tasks. Figure 4*B* illustrates that on average subjects followed the general trends predicted by the risk-neutral movement planner, and conditions of increased penalty and variability resulted in movement endpoints that are further from the cliff edge. On average, however, subjects moved closer to the cliff edge than predicted by the risk-neutral model. This is particularly evident in the WB task.

We performed a three-way repeated-measures ANOVA on movement endpoint data to examine the effects of risk condition and movement task. The levels were penalty (0 or −500

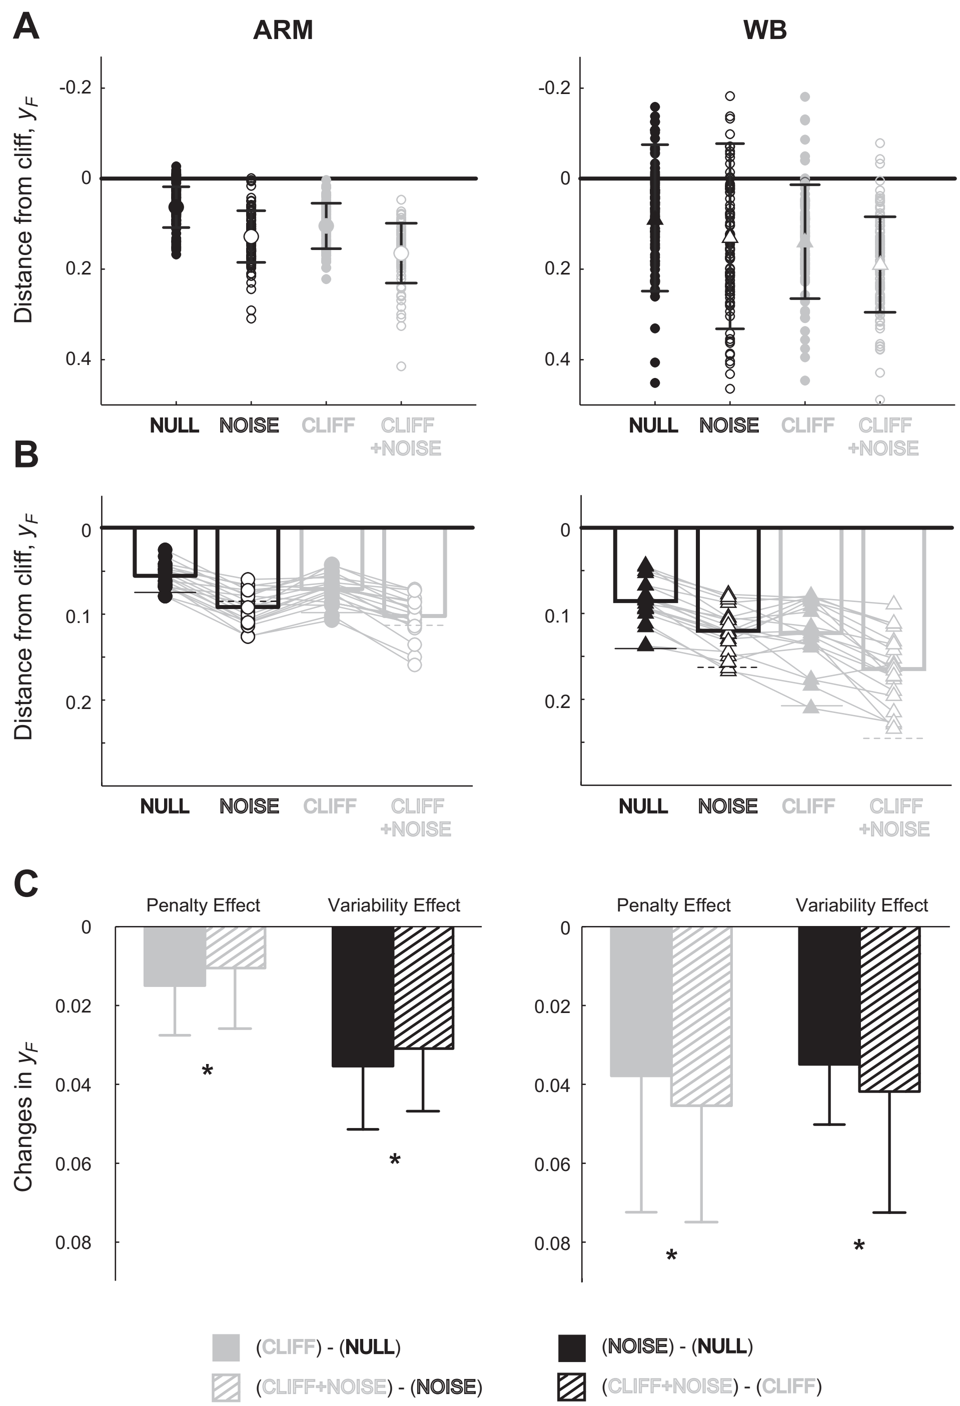

Fig. 4. Movement trends. *A*: movement endpoints, including mean and standard deviation, for S8 during the 4 conditions for ARM (*left*) and WB (*right*). Endpoints are expressed as distance from the cliff, y_F, and are normalized by the subject's cliff distance. The cliff edge is shown as a solid black line. *B*: mean movement endpoints normalized by cliff distance for all subjects during the 4 conditions for ARM (*left*) and WB (*right*). Mean subject endpoints for each condition are denoted by the filled or outlined bars, whereas mean endpoints predicted by the risk-neutral model for each condition are denoted by single horizontal lines. *C*: independent effects of penalty and variability on movement distance for ARM (*left*) and WB (*right*). Effect of penalty is determined by subtracting mean endpoints of NULL from CLIFF (gray shaded bar) as well as NOISE from CLIFF+NOISE (gray hatched bar). Effect of variability determined by subtracting mean endpoints of NULL from NOISE (black solid bar) as well as CLIFF from CLIFF+NOISE (black hatched bar). *$P < 0.05$, significant difference from zero.

points), variability (no added Gaussian noise and added Gaussian noise), and movement task (ARM and WB). We found independent effects of penalty, variability, and task ($P < 0.0001$), as well as a task × penalty interaction effect ($P < 0.002$). Thus the high cliff penalty prompted a greater change in movement endpoints for the WB task than for the ARM task. A subsequent two-way repeated-measures ANOVA of the risk condition at each level of movement task revealed independent effects of both penalty and variability for ARM and for WB (Fig. 4*C*), indicating that adding the high cliff penalty and adding cursor noise significantly affected the endpoint for both movement tasks (P values <0.005).

Risk-sensitivity. Mean risk-sensitivity values, calculated from *Eq. 5*, are shown in Fig. 5*A* for each subject, condition, and task. Group mean values are plotted in Fig. 5*B*.

Independent *t*-tests showed no significant difference between subjects' mean risk-sensitivity and 0% for the NOISE condition in the ARM task, whereas risk-sensitivity was greater than 0% in all other conditions in ARM and WB ($P < 0.05$). This is indicative of consistent risk-seeking behavior in both movement tasks. In the WB task, only one subject (S12) had a mean risk-sensitivity less than 0% in any condition (NOISE, CLIFF, CLIFF+NOISE). The same subject also had a mean risk-sensitivity less than 0% in all four ARM condi-

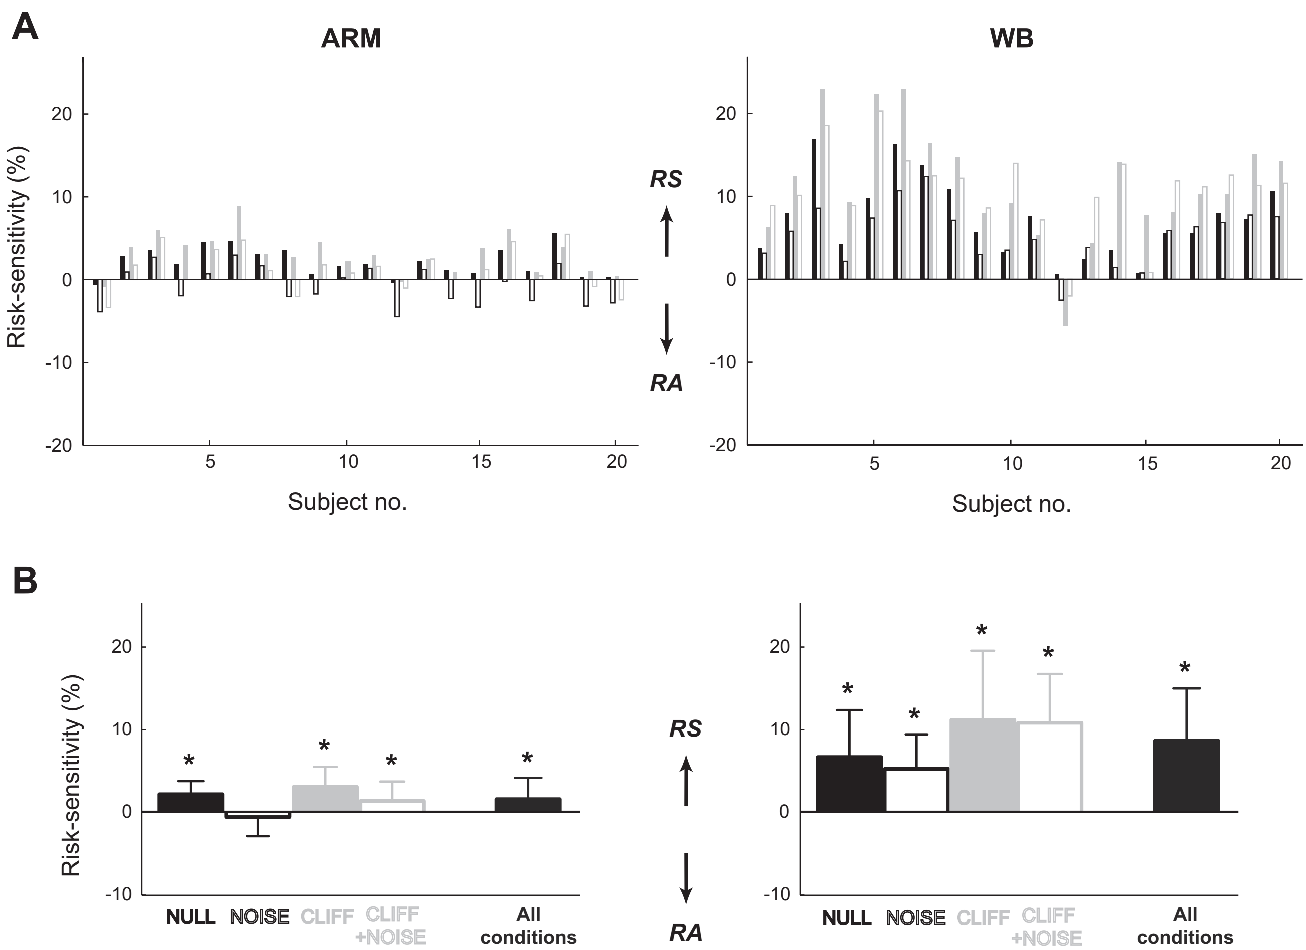

Fig. 5. Risk-sensitivity. *A*: mean risk-sensitivity values for individual subjects in each of the 4 conditions for ARM (*left*) and WB (*right*). A value <0% indicates that a subject did not move as close to the cliff as predicted (risk-averse, RA), whereas a value >0% indicates that a subject moved closer to the cliff than predicted (risk-seeking, RS). *B*: mean risk-sensitivity across subjects in each condition and across all conditions. $*P < 0.05$, significant difference from 0%.

tions, indicating relatively consistent risk-averse behavior for this subject. Only one subject (S1) demonstrated idiosyncratic risk preferences between movements in all conditions, with risk-averse behavior in ARM and risk-seeking behavior in WB. A paired *t*-test showed that risk-sensitivity values were significantly further from 0% in the WB task than in the ARM task for all four conditions (P values <0.002). Overall, subjects moved closer to the cliff in the WB task than the model predicted, and the discrepancy between actual endpoints and model-predicted endpoints was larger in the WB task than in the ARM task.

We turn to our second analysis to determine the consistency of risk-sensitivity across tasks for each subject. Figure 6 further illustrates the consistent direction of risk-sensitivity between movement tasks and the disparate degree of risk-sensitivity between movement tasks. Of the 80 available data points quantifying average risk-sensitivity (20 subjects × 4 conditions), 62 points were either risk-seeking in both ARM and WB or risk-averse in ARM and WB. This means that the direction of risk-sensitivity transferred across the two movements in 77.5% of all cases. A least-squares linear regression of group WB risk-sensitivity against ARM risk-sensitivity resulted in a regression slope of 7.2 ($R^2 = 0.30$, $F = 22.3$, $P < 0.0001$), confirming that subjects were more risk-seeking in WB. This finding held when we performed the same linear regression at the subject level; across subjects, the slopes of the regression line between conditions were significantly greater than unity ($P < 0.001$), with a mean (±SD) slope of 6.1 (4.1) and a mean (±SD) R^2 of 0.36 (0.29).

Our final analysis of risk-sensitivity sought to establish whether these deviations from 0% risk-sensitivity could be a manifestation of distorted utility or distorted probability weighting. We fit the weighting parameters α, β, and γ from *Eqs. 6a* and *6b* for each subject. Altering the utility and probability weighting functions shifts subjects' mean risk-sensitivity closer to 0% for both the ARM and WB movements, indicating that distortion of the α, β, or γ values could explain subject behavior during our experiment. In the ARM movement, our mean fit values were $\alpha = 1.13 \pm 0.17$, $\beta = 0.76 \pm 0.31$, and $\gamma = 1.13 \pm 0.22$. In the WB movement, our mean fit values were $\alpha = 1.42 \pm 0.27$, $\beta = 0.33 \pm 0.37$, and $\gamma = 1.22 \pm 0.17$. Across subjects, the three fit weighting parameters were significantly different from 1.0 in both movement tasks. These parameter fits corroborate our experimental observations of consistent risk-sensitivity across movements. Consistent with risk-seeking behavior, most subjects (17 of 20) overvalued the point rewards ($\alpha > 1.0$), undervalued penalties ($\beta < 1.0$), or overestimated their movement accuracy ($\gamma > 1.0$) in both movements. Only one subject (S12) had parameters that align with risk-averse behavior in both movements, and two subjects

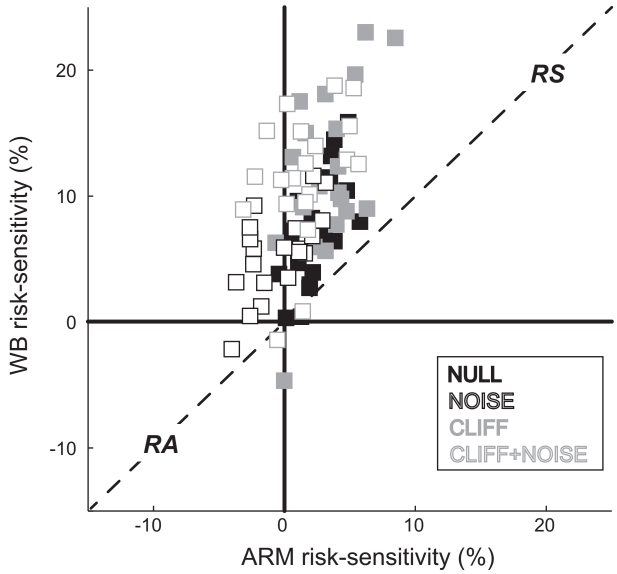

Fig. 6. Degree of risk-sensitivity. Risk-sensitivity for individual subjects in each condition, comparing ARM and WB. Data points that fall in the *top right* quadrant correspond to RS behavior. Data points that fall in the *bottom left* quadrant correspond to RA behavior. Unity is shown as a dashed black line. Least-squares linear regression for this comparison yields a slope of 7.2, confirming that the degree of risk-sensitivity is greater in the WB task (i.e., more risk-seeking in WB than ARM).

(S1 and S20) had parameters that showed idiosyncratic risk-sensitivity (risk-averse in ARM, risk-seeking in WB). For the two utility parameters α and β, there were significant differences between the ARM and WB tasks (α: $P = 0.0002$; β: $P = 0.0001$), indicating that distortions were larger in the WB movement. This also supports our behavioral findings of greater risk-seeking behavior in whole-body movements. There was not a significant difference in the variability parameter γ between ARM and WB ($P = 0.087$).

Effects of learning. Subjects did not appear to learn a movement strategy during the course of the experiment for either movement task. To determine whether subjects learned during the experiment, we examined both movement error and standard deviation of endpoints during each condition. For our purposes, movement error is equivalent to y_F, the distance between a subject's endpoint and the cliff edge. Most subjects did not exhibit a significant change in movement error during any given condition when the first 10 trials and the last 10 trials were compared (P values >0.05). The only exceptions to this were S2 (NULL: $P = 0.024$), S3 (CLIFF+NOISE: $P = 0.008$), S11 (CLIFF: $P = 0.005$), and S13 (CLIFF+NOISE: $P = 0.002$) in the ARM task, as well as S6 (CLIFF: $P = 0.038$, CLIFF+NOISE: $P = 0.025$), S8 (CLIFF+NOISE: $P = 0.037$), and S16 (NOISE: $P = 0.039$) in the WB task. There was no significant change in the standard deviation of movement endpoints between the first 10 trials and the last 10 trials (P values >0.05) in any condition and across all subjects.

Variability testing. One challenge in determining a risk-neutral movement endpoint is that movement variability changes as a function of time and distance. We had accounted for this changing variability in a separate experiment for each movement task, which allowed us to estimate subject-specific variability as a function of distance. For example, each subject's measured sensorimotor variability at the cliff distance, $\sigma_M(y_{cliff})$, and the estimated variability at the cliff distance, $\sigma_{0.66}$, is provided in Table 1. We determined these variability functions from the averages of each subject's PRE and POST variability measurements. However, POST variability data were not available for subjects S1–S4, since we did not instate POST testing until after these four subjects had completed testing. Furthermore, we did not use the PRE data for S12 WB or the POST data for S15 WB because these subjects did not follow Fitts' law in these cases. Specifically, variability decreased with movement distance in these two instances. Since this behavior was not consistent between the PRE and POST tests for these subjects, we attributed the results to external factors such as distracted attention or fatigue and questioned their validity as a true representation of the subject's sensorimotor variability.

The mean ($\pm$SD) values of the measured variability when subjects moved to the 66% target line were 0.53 (0.15) cm in ARM and 0.53 (0.14) cm in WB. Note that this average variability only includes S5–S20, since the 66% target line was not incorporated for the first four subjects. With the exception of S3, S4, S5, S6, and S17 in ARM, all estimated variabilities at the cliff distance were equal to or slightly less than the measured values. This indicates that the estimated functions $\sigma_M(y)$ were adequate representations of subject-specific sensorimotor variability. Estimating a lower variability than the measured values would produce model predictions that are slightly closer to the cliff. This means that the difference between actual endpoints and risk-neutral predicted endpoints would be smaller for subjects who move beyond the model predictions and larger for subjects who do not move as close to the model predictions. From *Eq. 5*, underestimating a subject's

Table 1. *Subject variability and probability distortion*

	ARM		WB	
Subject	$\sigma_M(y_{cliff})$, cm	$\sigma_{0.66}$, cm	$\sigma_M(y_{cliff})$, cm	$\sigma_{0.66}$, cm
S1	0.39	(0.42)*	0.36	(0.42)*
S2	0.56	(0.82)*	0.60	(0.60)*
S3	0.84	(0.82)*	0.51	(0.66)*
S4	0.50	(0.44)*	0.48	(0.48)*
S5	0.50	0.48	0.44	0.66
S6	0.55	0.53	0.57	0.60
S7	0.62	0.64	0.60	0.67
S8	0.46	0.46	0.68	0.69
S9	0.58	0.63	0.31	0.57
S10	0.67	0.73	0.70	0.73
S11	0.76	0.88	0.41	0.67
S12	0.61	0.61	0.32	0.52†
S13	0.46	0.54	0.22	0.25
S14	0.30	0.33	0.38	0.47
S15	0.42	0.42	0.36	0.38*
S16	0.59	0.62	0.33	0.33
S17	0.47	0.46	0.39	0.51
S18	0.41	0.44	0.34	0.38
S19	0.43	0.43	0.45	0.48
S20	0.26	0.28	0.59	0.62

Values are estimated sensorimotor variability at the cliff distance, $\sigma_M(y_{cliff})$, and measured variability at the cliff distance [averaged from variability testing before (PRE) and after (POST) the 4 risk conditions], $\sigma_{0.66}$, for each subject and movement task. Since we did not introduce the 66% target line or POST variability testing until after the first four subjects, $\sigma_{0.66}$ values given for S1–S4 are interpolated from variability at the 60% and 80% target lines during PRE testing and are shown in parentheses. *From PRE variability data only. †From POST variability data only.

sensorimotor variability would thus result in decreased risk-sensitivity values across conditions (portraying a risk-seeking individual as less risk-seeking, a risk-averse individual as more risk-averse, and a risk-neutral individual as risk-averse).

Generally, subjects tended to move slightly past the target line, overshooting by an average of 1.1% in the ARM movement and 4.8% in the WB movement. Interestingly, subjects were able to move nearly as precisely in the WB movement as in the ARM movement during this variability test. For example, the measured endpoint variability at the 0.66% target line, $\sigma_{0.66}$, was comparable between the two tasks (see Table 1); a paired *t*-test revealed that there was no significant difference in $\sigma_{0.66}$ between ARM and WB (P = 0.94, S5–S20). This was also true for the 40, 60, and 80% target lines, with no significant differences between ARM and WB measured endpoint variability (P values >0.3, S1–S20). Only for the 20% line was subject endpoint variability lower in the ARM task than in the WB task (P < 0.001).

Model adjustments. It is certainly possible our subject-specific variability functions did not accurately capture the subject's true endpoint variability and led to different estimates of risk-neutral movement behavior and thus different values of subject risk-sensitivity. We examined alternative variability functions as inputs to the SDT model and the resulting effects on risk-sensitivity. Considered alternatives included *1*) parameters from PRE testing, *2*) parameters from POST testing, *3*) a constant sensorimotor variability from PRE and POST testing, and *4*) a constant sensorimotor variability from each cliff condition. However, these adjustments did not yield any consistent decrease in risk-sensitivity for our subjects. Most often, these model adjustments would precipitate values that deviated even further from 0%. Overall, we feel confident that with our available data, we have presented a "best-case" scenario in comparing subject endpoints to risk-neutral model predictions.

DISCUSSION

This is the first study to assess risk-sensitivity in goal-directed whole-body movements and to compare risk-sensitivity across two dissimilar movements. We have shown that increasing risk in the form of point penalties and/or variability does affect movement endpoints in both arm-reaching and whole-body movements, and we present evidence of risk-sensitivity in whole-body movements. Overall, our findings demonstrate that subjects were generally risk-seeking in both movements. However, the degree of risk-sensitivity did not transfer between the two movements. In this section, we discuss each movement in turn, compare behavior in the two movements, and provide possible explanations for the differences between them.

In the arm-reaching task, movement endpoints are slightly closer to the cliff edge than predicted by the SDT model in three of the four risk environments, indicating that subjects are risk-seeking in these movements. This does not support previous findings of risk-neutral planning behavior in the hand/arm system under symmetric expected gain landscapes (Trommershäuser et al. 2003, 2008). Indeed, predicted endpoints more closely matched subject endpoints when we incorporated distortions in utility and probability weightings into the model that were distinctive of risk-seeking behavior. However, Wu et al. (2006) found that subjects performed suboptimally when pointing in an asymmetric expected gain landscape. In this case, the authors conjecture that subjects are not able to maximize expected gain in an asymmetric environment due to the increased complexity of the movement planning task. The expected gain landscape presented in our study is different than that of Wu et al. but is still inherently asymmetric, supporting their findings of suboptimal behavior.

In the whole-body leaning movement, increasing risk also significantly affected subjects' movement endpoints. A number of previous studies have shown an effect of implicit postural threat (i.e., an elevated support surface) on the control of COP (Adkin et al. 2000, 2002; Brown et al. 2006; Carpenter et al. 2001, 2006; Davis et al. 2009), but this is the first study to quantify the influence of explicit risk on goal-directed COP movements in a decision-making framework. Movement endpoints in this task are decidedly closer to the cliff than predicted by the model, indicating that subjects are not risk-neutral in whole-body leaning movements. Again, risk-seeking distortions in utility and probability generated model-predicted endpoints that more closely matched subjects' endpoints, further attesting to a general risk-seeking attitude during this movement. This evidence of risk-seeking behavior is unexpected, particularly in light of previous findings demonstrating that increased postural threat leads to more cautious whole-body movement control (Adkin et al. 2002).

Despite the distinct biomechanical differences between these two types of movement, we found that subjects adapt similar directionality in their risk-sensitivity. Recent studies have investigated whether risk-sensitivity transfers across other decision-making tasks and domains at the behavioral and neurobiological levels. Wu et al. (2009) observed dissimilar direction of risk-sensitivity in the same subjects in different decision making domains: a financial task and a movement task. However, when comparing tasks involving either food or financial rewards, Levy and Glimcher (2011) found that directionality of risk-sensitivity is correlated. Their corresponding neuroimaging results suggested that there may be overlapping neural substrates for risk-sensitivity across task domains. Likewise, our results suggest that there may be overlapping substrates for risk-sensitivity in movement control, but individuals can maintain different degrees of risk-sensitivity.

Our other main finding is that while the direction of risk-sensitivity was similar between movements, we did not see a similar degree of risk-sensitivity. Subjects were more risk-sensitive in the whole-body movement compared with arm-reaching. Thus an individual's movement decision making in one movement does not fully predict that individual's performance in another. Subjects tended to adopt a risk-seeking strategy during arm-reaching, but they were markedly more risk-seeking in the whole-body movement. Why doesn't the degree of risk-sensitivity directly transfer from one movement to another? Our model indicates that subjects could possess distorted utility and probability weighting functions and that these distortions differed between movements. In our experiment, such distortions could have arisen from *1*) an inaccurate estimation of sensorimotor variability, *2*) a subject not responding appropriately to the explicit point-based rewards/penalties, or *3*) our model of risk-neutral movement planning overlooks an underlying cost of motor control and is unable to predict subject behavior. We next address each of these possibilities in turn.

Risk-seeking behavior could result from an inappropriate estimation of task-related sensorimotor variability, by either the subject or the experimenter. A subject may internalize an estimation of their variability that causes overweighting large probabilities and underweighting small probabilities (distortion in probability). This means that subjects believe themselves to have a smaller endpoint variability than they actually do, which would most likely influence them to move closer to the cliff edge than predicted. Wu et al. (2009) observed a similar distortion pattern during a rapid pointing motor task, when the probability in question was simply subjects' own implicit sensorimotor uncertainty. We believe that inaccurate variability estimation is a manifestation of unfamiliarity with the motor task. Although forward leaning movements are relatively common in everyday tasks (such as when reaching for a cup in a high cabinet), such movements tend to involve slow, small leaning distances. However, the rapid, "out-and-back" goal-directed COP movements utilized in our experiment are difficult and are not often experienced on a daily basis. This could account for an inability to appropriately internalize one's sensorimotor variability within the duration of this experiment.

It is possible that our estimation of subjects' variability does not accurately reflect their true endpoint variability. Because our quantification of risk-sensitivity is dependent on this variability estimation, our results may be biased toward a consistent direction or degree of risk-sensitivity. We explored alternative measurements of variability to adjust our SDT model; however, these adjustments typically resulted in increased risk-sensitivity values (subjects were more risk-sensitive). We feel that we have presented a best-case scenario with our variability test data to compare subject endpoints with risk-neutral model predictions.

Alternatively, risk-seeking behavior could stem from overweighting point rewards and underweighting point penalties associated with the cliff paradigm (distortion in utility). However, we find this unlikely for two reasons. First, the point structure was the same in both the arm-reaching and whole-body tasks, so there is no obvious reason why the same subject would value these points differently between tasks. Even if a distorted utility function does exist, this distortion should remain consistent between tasks, which is not what we observe. Second, if anything, leaning forward closer to the limits of stability while standing should lead to an overweighting of the point penalties and result in more risk-averse behavior compared with the arm-reaching task. Leaning forward inherently increases the chance of a fall, thereby adding an implicit penalty over and above the explicit point penalties presented to the subject (Adkin et al. 2002). Individuals do indeed alter postural control strategies under greater implicit penalty, such as that imposed by standing on an elevated platform. These altered strategies, including decreased maximum reach, slight backward leaning, or reduced variability of postural sway, arguably align with risk-averse behavior and have been associated with both voluntary (Hauck et al. 2008) and involuntary leaning (Davis et al. 2009; Brown et al. 2007).

Given the dissimilar natures of arm-reaching and whole-body movements, it is certainly possible that these movements incorporate different control strategies or motor costs that are not included in our model. One potential candidate is an effort cost of muscle activation. In the arm-reaching movement, effort cost increases with movement distance. Adding an effort cost would lead to shorter movements, similar to the effect of increased variability. However, the relationship between COP movement and effort in this rapid out-and-back task has yet to be determined. In the WB task, the biomechanical properties of the foot/ankle are very different from the arm/hand and may contribute significantly to effort quantification or variability in the WB task. For instance, it could more desirable to move the COP closer to the balls of the feet because of more comfortable or familiar activation strategies in the foot/ankle, or due to some heightened control capabilities afforded by the COP's proximity to the toes. Theoretically, this could lead to a distortion in the underlying utility function, which would be different in the whole-body task than in arm-reaching. If moving greater distances required less effort or offered greater control, this would be manifested in the subject's behavior as an overvaluation of reward and undervaluation of penalty. Furthermore, subjects may be knowingly choosing a "satisficing" strategy over an optimal strategy (Simon 1956), and since the costs and familiarity of the whole-body movement are conceivably different than an arm-reach, this may also explain the larger deviations from model predictions seen in the WB task.

We used a CPT analysis to fit possible subject-specific distortions in the weighting of reward, penalty, and variability. The resulting trends in parameters support our findings of risk-seeking behavior in both movements and a greater degree of risk-seeking behavior in the whole-body task. Furthermore, the CPT analysis indicates that this risk-seeking behavior is not solely a result of variability distortions (and thus potential misestimates of variability by the subject). Rather, the CPT fits point toward a distortion in the subjective value of the gain landscape. Although the ability to verify the accuracy of this fitting procedure is always a concern, we hope to further address this in future studies.

In this experiment, we characterize movement in the arm-reaching and whole-body tasks with hand endpoint and COP endpoint, respectively. The COP is related to, but not equivalent to, the body's center of mass (COM). During upright postural adjustments, COP is the controlling variable, whereas COM is the controlled variable (Winter et al. 1998). The COP overshoots the vertical projection of the COM to keep the COM within the base of support, and the difference COP − COM is highly correlated with the negative acceleration of the COM. The difference between COP and COM is evident at the endpoint of a forward voluntary lean (Mancini et al. 2008) and becomes even greater during rapid out-and-back movements (Murnaghan et al. 2009), as is required in our study. Thus we are comparing incongruous measures between the two movement tasks (controlled variable in ARM and controlling variable in WB), and the COP (controlling variable in WB) is overshooting the COM (controlled variable in WB) at the movement endpoint. Initially, this may appear to explain the differences in degree of risk-seeking behavior observed between the ARM and WB tasks. However, we argue that using COP feedback in the WB task rather than COM does not change our interpretation of risk-sensitivity. In both movement tasks, we provide subjects with explicit visual feedback of the variable they are being asked to control (hand position in ARM, COP position in WB). The fact that COM will always lag behind COP at the movement endpoint does not ensure that subjects will move their COP beyond the target provided in the WB task. Rather, they likely control their movements such that the COP lands on or near the target, and the COM will be influenced by these dynamics accordingly. We see no reason why subjects would control their COM when specifically asked to control their COP and given the necessary feedback to do so. Indeed, a previous study from our laboratory confirms that subjects can control their COP in a goal-directed manner when instructed to do so, even when the target changes location mid-movement (Huang and Ahmed 2011). Even if they controlled their COM at the onset of the present experiment, one would expect them to adjust any unexpected endpoint behavior based on the scoring feedback provided. However, we generally do not observe an effect of learning or altered movement strategies throughout the duration of a trial set. Furthermore, all subjects were able to move to various distances on the "safe" side of the cliff, as evidenced by performance during the PRE and POST variability testing. Nevertheless, many subjects would express verbal disappointment if they did not earn a score in the 80s or 90s, though they might repeatedly move past the edge of the cliff. We therefore do not attribute greater risk-seeking behavior in WB to an impaired ability to control COP over COM, but instead consider risk-sensitive behavior as a manifestation of the aforementioned possibilities (inappropriate variability estimation, inappropriate reward/penalty weighting, or an unaccounted cost).

It should also be noted that we did not scale the Gaussian cursor noise to the cliff penalty, so our various manipulations of risk (increasing variability and/or penalty) may not be equivalent. This could explain why we observe greater risk-sensitivity under conditions of increased penalty (CLIFF over NULL and CLIFF+NOISE over NOISE) than under conditions of increased variability (NOISE over NULL and CLIFF+NOISE over CLIFF). We expect that adding cursor noise with an even larger amount of standard deviation would still cause risk-seeking behavior in both movement tasks. This behavior does not align with that observed by Nagengast et al. (2010), where subjects acted risk-averse in the presence of cursor noise during a goal-directed arm-reaching movement in the horizontal plane. Nonetheless, these findings highlight the potential sensitivity of results to experimental context and the importance of controlling for context and inter-individual differences in this and future studies.

Our findings have important implications for quantitative descriptions of decision making to generalize across movements and, ultimately, across decision-making contexts. Further research is required to determine whether and to what extent risk-sensitivity transfers across these contexts, the neural structures governing such mechanisms, and under what conditions we may observe transfer across decision-making domains.

ACKNOWLEDGMENTS

We thank Prof. Rodger Kram for supplying the force plate used in this experiment.

GRANTS

This work was supported by Defense Advanced Research Projects Agency Young Faculty Award DARPA YFA D12AP00253 and National Science Foundation Grants SES 1230933 and CMMI 1200830.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

M.K.O. and A.A.A. conception and design of research; M.K.O. performed experiments; M.K.O. analyzed data; M.K.O. and A.A.A. interpreted results of experiments; M.K.O. prepared figures; M.K.O. drafted manuscript; M.K.O. and A.A.A. edited and revised manuscript; M.K.O. and A.A.A. approved final version of manuscript.

REFERENCES

Adkin AL, Frank JS, Carpenter MG, Peysar GW. Postural control is scaled to level of postural threat. *Gait Posture* 12: 87–93, 2000.

Adkin AL, Frank JS, Carpenter MG, Peysar GW. Fear of falling modifies anticipatory postural control. *Exp Brain Res* 143: 160–170, 2002.

Braun DA, Nagengast AJ, Wolpert DM. Risk-sensitivity in sensorimotor control. *Front Hum Neurosci* 5: 1–10, 2011.

Brown LA, Doan JB, Whishaw IQ, Suchowersky O. Parkinsonian deficits in context-dependent regulation of standing postural control. *Neurosci Lett* 148: 292–297, 2007.

Brown LA, Polych MA, Doan JB. The effect of anxiety on the regulation of upright standing among younger and older adults. *Gait Posture* 24: 397–405, 2006.

Carpenter MG, Adkin AL, Brawley LR, Frank JS. Postural, physiological and psychological reactions to challenging balance: does age make a difference? *Age Ageing* 35: 298–303, 2006.

Carpenter MG, Frank JS, Silcher C, Peysar G. The influence of postural threat on the control of upright stance. *Exp Brain Res* 138: 210–218, 2001.

Davis JR, Campbell AD, Adkin AL, Carpenter MG. The relationship between fear of falling and human postural control. *Gait Posture* 29: 275–279, 2009.

Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci* 9: 292–303, 2008.

Faisal AA, Wolpert DM. Near optimal combination of sensory and motor uncertainty in time during a naturalistic perception-action task. *J Neurophysiol* 101: 1901–1912, 2009.

Hauck LJ, Carpenter MG, Frank JS. Task-specific measures of balance efficacy, anxiety, and stability and their relationship to clinical balance performance. *Gait Posture* 27: 676–682, 2008.

Huang HJ, Ahmed AA. Tradeoff between stability and maneuverability during whole-body movements. *PLoS One* 6: 1–10, 2011.

Levy DJ, Glimcher PW. Comparing apples and oranges: using reward-specific and reward-general subjective value representation in the brain. *J Neurosci* 31: 14693–14707, 2011.

Mancini M, Rocchi L, Horak FB, Chiari L. Effects of Parkinson's disease and levodopa on functional limits of stability. *Clin Biomech* 23: 450–458, 2008.

Markowitz H. Portfolio selection. *J Finance* 7: 77–91, 1952.

Murnaghan CD, Elston B, Mackey DC, Robinovitch SN. Modeling of postural stability borders during heel-toe rocking. *Gait Posture* 30: 161–167, 2009.

Nagengast AJ, Braun DA, Wolpert DM. Risk-sensitive optimal feedback control accounts for sensorimotor behavior under uncertainty. *PLoS Comput Biol* 6: e1000857, 2010.

Simon HA. Rational choice and the structure of the environment. *Psychol Rev* 63: 129–138, 1956.

Trommershäuser J, Maloney LT, Landy MS. The expected utility of movement. In: *Neuromechanics*, edited by Glimcher PW, Camerer C, Fehr E, Poldrack R. New York: Academic, 2008, p. 95–111.

Trommershäuser J, Maloney LT, Landy MS. Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A Opt Image Sci Vis* 20: 1419–1433, 2003.

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

van Beers RJ, Baraduc P, Wolpert DM. Role of uncertainty in sensorimotor control. *Philos Trans R Soc Lond B Biol Sci* 357: 1137–1145, 2002.

Winter DA, Patla AE, Prince F, Ishac M, Gielo-Perczak K. Stiffness control of balance in quiet standing. *J Neurophysiol* 80: 1211–1221, 1998.

Wu SW, Delgado MR, Maloney LT. Economic decision-making compared with an equivalent motor task. *Proc Natl Acad Sci USA* 106: 6088–6093, 2009.

Wu SW, Trommershäuser J, Maloney LT, Landy MS. Limits to human movement planning in tasks with asymmetric gain landscapes. *J Vis* 6: 53–63, 2006.