

Finite-time thermodynamics of an instrumented drinking bird toy

Ralph Lorenz^{a)}

Lunar and Planetary Lab, University of Arizona, Tucson, Arizona 85721

(Received 8 November 2005; accepted 3 March 2006)

The motion and temperature of a drinking bird toy is monitored by a variety of instruments to determine the quantitative history of its motion over long times and to determine the thermodynamic and mechanical constraints on its performance. © 2006 American Association of Physics Teachers.
[DOI: 10.1119/1.2190688]

I. INTRODUCTION

The “drinking bird” (also called “dipping bird,” “dippy bird,” and other names) is an appealing toy that demonstrates some interesting thermodynamic principles by maintaining periodic motion for days without an obvious supply of energy. In essence, the toy is a pair of glass bulbs linked by a pipe and partially filled with a volatile liquid. The glass assembly is mounted on a pivot, allowing it to rotate about a horizontal axis. The upper bulb (the head) has an attached beak, moving the center of mass slightly forward, and is covered in a thin felt that retains water. When this felt is made wet in a dry environment, the water slowly evaporates, cooling the upper bulb. The lower temperature reduces the vapor pressure of the volatile fluid in the head relative to the lower bulb, and the liquid rises up the pipe (the lower end of which is mounted so that it is immersed in the liquid). As the liquid rises into the head, the center of mass of the system rises to the point that the assembly tips forward, dunking the beak into a glass of water to keep it moist. When the assembly is tipped forward, the lower end of the pipe gains access to the vapor in the lower bulb (the abdomen), which then bubbles up the pipe, allowing the liquid in the head to flow back into the lower bulb. Like a pendulum, the bird swings back and forth before the cycle begins again.

In dynamical terms, the system is a damped compound pendulum whose mass varies with time due to the slow upward migration and sudden draining of an internal liquid. In thermodynamic terms, it is a heat engine that uses the evaporation of water to create a temperature difference and heat flow. (The origin of the motive power is the disequilibrium of liquid water in an unsaturated environment—in essence, a concentration potential such as that in a chemical battery).

Among previous work on the topic (see the comprehensive paper by Guemez *et al.*¹ and references therein), there are only a few papers that quantitatively study either the motion of the bird or its thermodynamics² as an engine.

A detailed study by Murrow³ documents experiments on a drinking bird configured to raise water and suggests the application of scaled-up engines of this sort for irrigation of crops near the Nile, but is difficult to obtain. A short article by Rathjen⁴ shows how a gear and ratchet arrangement allows a dipping bird to produce work by raising a small mass (a paperclip), with about 1 μW of useful work being generated.

A numerical model of the dynamics of the dipping bird¹ includes the parameterization of the water evaporation rate (essentially, a measure of the heat supplied to the engine) as a function of relative humidity. The period of the bird’s cycle was measured and reproduced by adjusting the model parameters. Although the model can describe the instantaneous position of the bird, this was not measured.

In this paper, I explore the behavior of the dipping bird (see Fig. 1). The instrumentation used allows the history of the angular position of the bird and the evolution of the dipping interval to be conveniently monitored over long times. Furthermore, the temperature evolution of the bulb surface is monitored to gain a sense of the heat transfers in the system. Finally, the ability of the device to develop electrical energy is shown. These data lead to a discussion of the factors behind the device’s low efficiency as an engine.

II. EQUIPMENT AND RESULTS

We used a small classic dipping bird (Edmund Scientific Part No. 3053617, manufactured by Asaro), which is about 15 cm long with bulbs about 26 mm in diameter. This bird is similar to those used in other studies,^{2–4} although larger than the perched variant used in Ref. 1. The bird was partly filled with dyed methylene chloride (dichloromethane).

A strong samarium-cobalt magnet (6 mm thick by 1 cm diameter) was attached to the bottom of the lower bulb using silicone adhesive, with the magnetic polar axis parallel to the bird’s pipe. Adhesive putty was applied to the hat of the bird to offset the magnet’s mass so that the bird still pivoted correctly.

A small coil (inside diameter 6 mm, outside diameter 17 mm, height 6 mm, resistance 410 Ω , and inductance 6.1 H) was placed beneath the bird’s lower bulb (see Fig. 2) and secured in place with adhesive putty. The leads from the coil were wired across a 330 Ω resistor.

A small one-axis magnetometer sensor (Speake FGM-1) was mounted about 20 cm from the base of the bird and secured with adhesive putty. The curved field lines from the magnet generate a field along the sensing axis of the device which varies monotonically with the angle of the bird. (The field also varies with relative location, but the location was held fixed.) The sensor output is a 5 V square wave, whose period is proportional to the field. By holding the bird at various angles, a calibration curve (fit adequately by a cubic polynomial—see Fig. 3) was constructed showing the monotonic relation of the output frequency to the bird angle. Fine (0.12 mm diameter) chromel and alumel wires were joined (forming *K*-type thermocouples) with the junctions mounted by a tiny smear of glue on the felt on the bird’s head and its abdomen.

The coil and thermocouple voltages and the magnetometer frequency were recorded by an Iotech PersonalDaq55 data acquisition unit, attached via a USB interface to a personal computer. This unit logged the sensor output at 0.515 s per sample. Some runs monitored only one sensor, at a higher



Fig. 1. Experimental setup. A magnet is placed on the base of the bird, just above the coil. Just behind the data logging unit on the right is a small cylindrical magnetometer mounted on putty. The voltmeter in the foreground was used to verify the thermistor excitation voltage. On the lower left is a handheld weather station for recording the temperature and humidity and a protractor mount for calibrating the magnetometer.

rate (6 Hz). The unit also corrects the thermocouple readings for the temperature of the unit itself, that is, cold junction compensation.

The tilt angle during one cycle is shown in Fig. 4. The bird slowly tips over, accelerating until it hits 40° and the pipe empties. It then swings back and forth, with a decreasing amplitude until the cycle repeats. As the amplitude of the oscillations decreases, the mean position also decreases due to the change of the center of mass as the liquid rises up the pipe. The fact that successive swings are not always smaller in the plot in Fig. 4 is probably an artifact of the low sampling rate, so that the actual peaks are not captured. The observed history appears very similar to that predicted by the model.¹

Figure 5 shows the coil voltage as a function of time. The time dependence is more difficult to interpret. Immediately after “drinking,” when the fluid drains from the head to the




Fig. 2. Closeup of the base of the bird. The magnet, coil, and the bird are held in place by adhesive putty. Fine wires from the thermistor can just be seen.

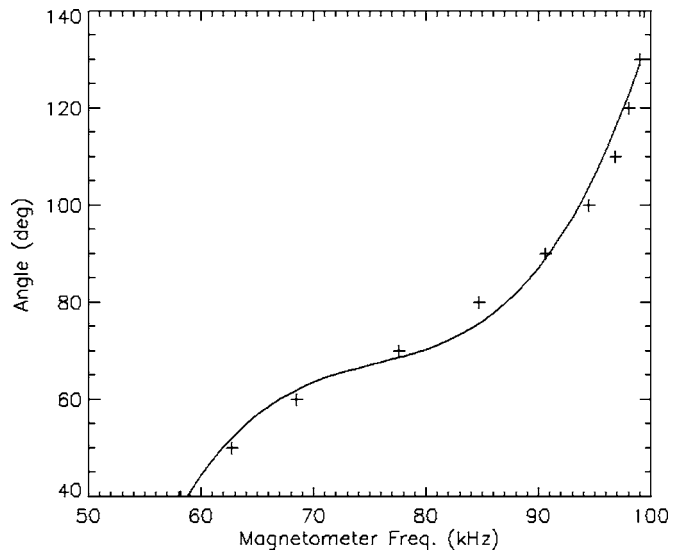


Fig. 3. Magnetometer output pulse frequency as a function of the tilt angle of the bird. A third-order polynomial provides an adequate fit to the data recovering the tilt angle to an accuracy of better than 5° .

abdomen, the bird acts as a pendulum with a near vertical central position near the coil. As the magnet sweeps past the coil (twice each swing), it generates a “heartbeat” pulse, shaped like an M because the field lines cross the coil in one direction, then the other, then back to the first. Later in the bird’s cycle, some liquid has ascended the pipe and the mean position is displaced. Despite its complexity, the time dependence in Fig. 5 is very reproducible from cycle to cycle. Because the large pulse developed at the center of the swing




Fig. 4. The angle of the dipping bird in one cycle as determined from the magnetometer frequency. An angle of 90° corresponds to the pipe being vertical. The bird slowly tips over, accelerating until it hits 40° and the pipe empties. It then swings back and forth with a decreasing amplitude until the cycle repeats. As the oscillations decrease, the mean position also decreases due to the change of the center of mass as the liquid rises up the pipe. The fact that successive swings are not always smaller is probably an artifact of the low sampling rate, so that the actual peaks are not captured.




Fig. 5. Coil voltage record (the voltage is zero for the other half of the cycle). The rapid motion begins symmetrically across the coil, leading to the M-shaped pulses at the beginning; later motions are centered around a tilted position such that the pulses become more sinusoidal. The cycle is very reproducible and the dotted line superimposes a record from a cycle 1 h later.

occurs only a fraction of a second after the bird finishes drinking, it is a useful marker for defining the start of a new cycle and defines the periods reported here.

The system generates a voltage across a load resistance (330 Ω). The peak instantaneous voltage is around 30 mV and thus the peak power is $V^2/R=0.032/330\approx 3\ \mu\text{W}$. If we average over the entire cycle (including the roughly 50% of the time that the bird is essentially motionless), we obtain an average power output of 0.27 μW . The load resistance was chosen to be roughly the same as the coil resistance. Although using a larger load resistance (or leaving the coil open circuit) would give a slightly higher measurement voltage, less of the power developed in the coil would be dissipated in the load. Conversely, using a lower load resistance would drop the voltage substantially, making it more difficult to measure. As is known from elementary circuit theory, the load power dissipation is maximized when the load resistance equals the internal resistance of the source.

Figure 6 shows the time dependence of the head and abdomen thermocouples in relation to the position data and coil voltage. It is seen that there is a steady-state difference of about 3 $^\circ\text{C}$ between the head and abdomen (as there must be for heat to enter the abdomen and be transported to the head). The temperature of both reservoirs varies by 0.1 K over each cycle. The thermodynamic significance of these measurements will be discussed in Sec. IV.

Both reservoirs, and especially the upper (cold) one, see rapid (≈ 1 s, 0.3 $^\circ\text{C}$) variations in temperature. These variations correlate with the angular motion of the bird, indicating heat transport from the air, and evaporation enhancement due to the relative wind caused by swinging.

Although thermocouples were used for the principal results reported here, an initial experiment used a thermistor. A small glass-encapsulated thermistor (10 k Ω at 20 $^\circ\text{C}$: Electronic Goldmine Part No. G1929) was wired as a potential divider with a 10 k Ω resistor and the pair was driven by a 1.5 V D cell. A low excitation voltage was used to minimize self-heating. The thermistor—the size of a grain of rice—was connected to the power source and datalogger with 38




Fig. 6. Record of three cycles. The top panel shows the angle measured by the magnetometer. After tilting over to $\approx 40^\circ$ where the bird drinks, it oscillates with an $\approx 30^\circ$ amplitude about the vertical, decreasing, and centering on a progressively more tilted position. The second panel shows the coil voltage, which is related to the speed through the vertical position (notice that the low sample rate gives a misleading one-sided view of the instantaneous voltages, which are seen better in Fig. 5). The third panel shows the temperature of the felt on the head: Pulses of warming occur during rapid motion as the thermocouple is warmed by the air, but immediately thereafter strong cooling occurs as the evaporative chill from nearby felt is conducted to the sensor. After cooling to a minimum as the motion stops, the temperature rises as the (warm) liquid rises into the head. The rise at 3475 s may be due to the beak being immersed in the liquid. The fourth panel shows the temperature of the lower bulb.

AWG enameled copper wire (magnet wire) which was sufficiently thin not to apply torque to the bird. Although requiring less sophisticated signal processing, the response time of the thermistor was too slow to fully capture the changes in temperature during the swings (see Fig. 7). However, if fine thermocouples are not available, thermistors permit the demonstration of some of the processes. (We could consider many alternatives or augmentations to all of the equipment used here, both in terms of the data acquisition hardware used and the sensing techniques applied.)

III. DIPPING BIRD PERIOD

Because the mechanical energy per cycle of operation is fixed, it follows that the output power of the engine can be expressed as the rate at which cycles occur. Thus, the period of the bird is related to the evaporation rate of water: this rate depends on the ambient humidity and temperature. An interesting quantitative dependence of the period on the humidity has been noted,¹ in particular, the evaporation rate switches below a relative humidity H of around 75% from being proportional to $(100-H)^{1.24}$ to a regime where it varies as $(100-H)^{1.82}$.

Murrow³ explored the dependence of the period on air-flow. As is well known to anyone who has to dried laundry outside, even a gentle wind can enhance evaporation significantly. The dipping period of Murrow's device (at a relative humidity of 64%) decreased by a factor of 6 by changing zero wind conditions to a flow of only 1 m/s, induced by a




Fig. 7. A similar record as in Fig. 6, but here the temperature been measured with a thermistor. The peaks and troughs are somewhat muted because the thermistor has a much larger thermal mass and thus a slower response than the thermocouple. The fact that the cooling is modulated by the swing is still evident. (The one-sided coil voltage record is again an artifact that would be eliminated by faster sampling.)

small fan. Increasing the wind to 2.5 m/s brought only a modest decrease in period (eight times smaller than for zero wind). Although we applied no wind, the more rapid swings with an amplitude of around 1 radian and a period of 1 s induce a head speed of around 0.5 m/s, and the temperature record (Figs. 6 and 7) shows the effect on cooling.

In the experiments reported here, the period changes rather significantly with a modest change in the temperature, with the relative humidity remaining essentially constant at around 35%. The period can be measured in a variety of ways. I used the coil voltage history. A program in IDL (Interactive Data Language by Research Systems) read the text file produced by the data logging program and searched for the time when the voltage exceeded a threshold (0.001 V). This value defines the start of a cycle. It then searched for the time when ten successive samples differ by less than another threshold (5 μ V), judging the motions to have damped. The program then searches for the next 0.001 V threshold crossing, and takes the difference in time between the two thresholds to be the cycle period. The method was verified by manually inspecting the record to measure the first 100 periods; the periods agreed in over 95% of cases to within 1 s.

As seen in Fig. 8, there are more or less random cycle-to-cycle variations in the period, with a standard deviation of about 2 s or 5%. These variations might depend on the specific (and possibly nondeterministic) details of the wetting process, together with possible air temperature fluctuations or air movements affecting the evaporation rate. (Murray employed a surfactant to eliminate surface tension effects in his experiments.) The variation of the period means that while the relative time spent in a given position and the average cycle period may be calculated knowing the ambient conditions, the exact position and the time of dipping cannot be determined *a priori* in advance. Thus, the dipping bird could be used as a thermally driven physical system that




Fig. 8. Periods of 100 successive cycles. The square points correspond to a period calculated automatically as described in the text. The triangles correspond to the same cycles measured from the voltage record by hand. The good agreement suggests that the cycle-to-cycle variations in period are real.

generates unpredictable seeds for random number generation, much like the Lavarand experiment using lava lamps.⁵

Over longer periods, the period shows significant variations (see Fig. 9) that are due to changes in the meteorological conditions affecting evaporation. In this indoor experiment, the relative humidity remained nearly constant over a 24 h interval, while the temperature varied from 23 to 26 °C. Even for this small change in temperature, the period varied by almost a factor of 2. The dependence of the




Fig. 9. Cycle periods for 24 h of operation (beginning at midnight). The ambient humidity was between 32 and 35%, with air temperature varying between 23 and 26 °C. Periods were measured automatically (points) as described in the text. The rapid changes in the period may be related to the air conditioning in the room. The solid line shows that in this near-constant humidity environment, the period correlates with room temperature, $P = 540 - 20T$, where P is in seconds and T is in Celsius.

period on room temperature is a linear expression of the form $P=540-20T$, where P is in seconds and the temperature T is in Celsius.

IV. FINITE-TIME PERFORMANCE OF THE DIPPING BIRD

Ng and Ng² calculated the efficiency of a dipping bird by dividing the work per cycle by the latent heat of the water evaporated. They compared this efficiency (450 cycles of ≈ 2.3 mJ, during which time 2 ml of water were evaporated, giving an efficiency of 0.02%) to the Carnot efficiency. The Carnot efficiency was evaluated using the ambient temperature and the wet bulb temperature as hot and cold limits, respectively (21.1 °C and 17.8 °C), giving a Carnot efficiency of 1.1%. As is well known, a practical engine can attain the Carnot limit only by operating infinitely slowly.

Murrow's investigation focused only on maximizing the work of the engine, not its efficiency. His bird, equipped with a scoop, could raise water by a few cm, and yielded about 1 μ W of power at a relative humidity of 64% with no wind. However, the water consumption rate was not recorded.

The present experiments allow a more detailed investigation into the engine's losses. Over the 24 h run, the glass into which the bird dipped lost 1 g/h of water. Most of the water was evaporated on the head of the bird. An identical glass placed adjacent to the dipping bird lost ten times less water from its 44 mm square surface than the one supplying the bird. (This loss suggests that the drinking bird engines could be used as room humidifiers.) Sometimes a drop of water appeared at the base of the bird, apparently having dribbled down from the felt. If this 5 mm drop was always present and has the same rate of evaporation per unit area as the control glass, it would only add a negligible (<1%) water loss, in essence a "fuel leak." The effective evaporation rate due to the bird's operation is ≈ 0.9 g/h, corresponding to a heat draw from the engine of about 0.55 W, or for a 40 s period, about 22 J per cycle.

The heat actually exploited by the engine corresponds to the latent heat needed to vaporize the 2 ml of vapor needed to displace the liquid. The pressure inside the glassware is not known, but adopting the vapor pressure of dichloromethane at 25 °C of 430 Torr $\approx 5.6 \times 10^4$ Pa, the 2 ml of vapor corresponds to about 3.5 mg of fluid. The latent heat of vaporization is 3.9×10^5 J/kg, and thus the heat actually processed in the cycle is ≈ 1.3 J. For a period of typically 40 s, this heat transport per cycle corresponds to a heat flow of 34 mW, or only 6% of the water evaporation rate.

Some of the heat flow induced by the evaporation will be lost by diffusion (for example, to the surrounding air) but this loss is probably minor. More important, as noted,² is that the liquid is warmed and cooled during each cycle, but the heat used in this way does not contribute to the operation of the engine. This point is important; although the motion of the liquid is the visible cause of the engine's motion, it is the invisible dichloromethane vapor that is the working fluid in the thermodynamic cycle. Assuming that the ≈ 2 mL of moving liquid is raised and lowered in temperature by the 0.1 °C observed by our thermocouple measurements, the heat wasted is $(0.002 \times 1.3 \times 0.1 \times 900) = 0.23$ J. Thus, the thermal inertia losses of heating and cooling components of the engine (both bulbs and the fluid) account for about as much heat as is actually used in the thermodynamic cycle.

However, if some fraction of the liquid is cycled by the steady-state difference between hot and cold bulb walls (3 °C), then this loss may be the most significant term, some 7 J per cycle, and the total heat flow through the system is ≈ 0.2 W. The other 0.35 W is presumably lost to the surroundings.

The temperature limits of the Carnot cycle are narrower than the limits given by the hot and cold reservoirs of the system, in which the engine operates because temperature gradients (with irreversible losses) must exist to draw heat into the engine and expel heat from it. The paradigm of a perfect reversible engine, connected to hot and cold reservoirs by thermal resistances, is known as the "endoreversible engine."

We see from Fig. 6 that the hot and cold ends of the dipping bird are about 3 K apart. These operating temperatures of 21.3 °C and 18.4 °C are in contrast to the ambient temperature of 24 °C and relative humidity of 31%. These conditions correspond at our altitude of 700 m to a dew point of 5 °C and a wet bulb temperature of 14 °C. Thus, an ideal Carnot efficiency² of $\approx 9/297 = 3\%$ would apply. (The experiments in Arizona were conducted in lower humidity than is typical.)^{1,2} The finite evaporation rate and heat transfer into the engine are such as to restrict the engine's efficiency to $\approx 3/294 \approx 1\%$.

The usual conductive losses in an endoreversible engine are not especially large. The bulbs have surface areas A of approximately 0.002 m² and the thickness of the glass is $t \approx 2$ mm. If we take the thermal conductivity k as unity (typical of glasses) in SI units and note that the heat flow F internal to the engine (liquid warming plus evaporation) is ≈ 0.2 W, we find that the temperature drop across the glass is only $\Delta T = Ft/kA \approx 0.2$ K. Poor conduction in the engine walls is thus not a major limiting factor in engine performance—the external vapor transport to and from the walls is more significant.

As an aside, we can demonstrate the consistency of the glass thickness that we assumed by observation. By watching closely, we observe that as the bird swings back and forth, the liquid column rises up the pipe in steps, suggesting that the vapor temperature inside the bulb must respond to the swing-modulated cooling, and thus the bulb walls must have a thermal time constant of around 1 s or less. Typical borosilicate glass has $\rho = 2200$ kg/m³, $k = 1.1$ W/mK, and heat capacity $c_p \approx 800$ J/kg K. The thermal diffusivity is thus $\kappa = k/\rho c_p \approx 6 \times 10^{-7}$ m²/s. If the glass on the bulbs is about 2 mm thick as suggested, we calculate a thermal time constant of $t^2/\kappa \approx 1$ s.

If we apply the 1% internal Carnot efficiency to the heat flow in the vapor cycle of 34 mW, we obtain an output of 340 μ W. Mechanically, each cycle of the engine corresponds to the raising of about 2 ml of liquid ($\rho \approx 1.3$ g/cm³) by about 10 cm and, hence, a work of about 2 mJ. Over a 40 s period, this work corresponds to a mechanical power of ≈ 50 μ W and a total system efficiency of $\approx 0.01\%$, comparable to the results in Ref. 2. (The difference between 50 μ W and the theoretical 340 μ W is not understood, and may be due to inhomogeneities in the temperature and heat flow in the bulb, which have led to underestimates of the losses.) However, even much of this energy is discarded during the operation of the engine. In the process of slowly leaning to drink, about $(1 - \cos 45^\circ)/2 \approx 15\%$ of the energy is lost compared to an ideal horizontal arrangement of the full and

empty bulbs (although some of this energy may be captured in some mechanical arrangements).^{3,4} Then, as the liquid drains down to the abdomen, nearly $\cos 45^\circ \approx 70\%$ is lost, because the flow of liquid (somewhat impeded by viscosity and the upward movement of bubbles) contributes little to the motion of the bird as a whole. Only the remaining $\approx 15\%$ is manifested as kinetic energy in the swinging of the bird (energy which might be possible to extract by means of a coil-magnet dynamo).

In our setup, in which the coil did not noticeably damp out the pendulum motion, only $0.27 \mu\text{W}$ was extracted; whereas our analysis suggests perhaps $7 \mu\text{W}$ might be available. We could use a lower-impedance coil to draw more power from the bird's swing, which would act as a better dynamo than our setup, which was chosen for measurement capability rather than energy conversion. However, a more effective dynamo would act as a magnetic brake and damp the swing more. Because the swinging is responsible for enhancing evaporation, improving the energy extraction from the bird would ultimately yield diminishing returns. Presumably, this motion-dependent evaporation rate is responsible for the regime change in the period versus humidity observed in Ref. 1.

V. CONCLUSIONS

Modest instrumentation has permitted the quantification of the behavior of a drinking bird toy and shows the variation of its period over a 24 h interval. In a domestic indoor setting, with fairly constant relative humidity and nominally zero windspeed, the period varies by almost a factor of 2, follow-

ing an external temperature change of only 3°C . A non-deterministic variation of the period of the order of 5% was measured.

Temperature measurements indicate the role of swinging in enhancing the evaporative cooling of the head. Although not all inefficiencies were identified, substantial thermodynamic losses are due to diffusive losses and the thermal inertia of temperature-cycled engine components other than the working fluid, namely, the liquid and the walls. Conductive losses in the walls are small.

Mechanically, most of the available work is lost as the liquid drains to the lower bulb. About 15% is manifested in the dipping action and about 15% in the swinging motion. Although some modest additions could improve the energy extraction from the dipping bird, as an engine it is far from ideal. However, it was never conceived as an energy conversion device, and its thermodynamic inefficiencies only make its performance as an exquisitely simple toy all the more impressive.

^{a)}Electronic mail: rlorenz@lpl.arizona.edu

¹J. Guemez, R. Valiente, C. Fiolhais, and M. Fiolhais, "Experiments with the drinking bird," *Am. J. Phys.* **71**(12), 1257–1263 (2003).

²L. M. Ng and Y. S. Ng, "The thermodynamics of the drinking bird toy," *Phys. Educ.* **28**, 320–324 (1993).

³R. B. Murrow, "A simple heat engine of possible utility in primitive environments," Report No. P-3367, RAND Corporation (1966).

⁴D. Rathjen, "Duckpower," *Exploring* **7** (Winter 1994).

⁵I. Peterson, "Lava lamp randomness," *Sci. News* (Washington, D. C.) **152**, 92 (1997).

LEARNING TO ACCEPT LIMITATIONS

Most people find it hard to believe that our minds impose fundamental limitations on our ability to explain the behavior of the world around us. We have no difficulty accepting the fact that our bodies impose limitations on our physical abilities. The speeds at which we run, the heights to which we jump, and the weights we can lift are ultimately limited by our physiology. None of us questions or complains about these limits. Although great athletes seem to be constantly extending them farther than we thought possible, the extensions become smaller and smaller and are increasingly difficult to achieve. Similarly, there are limits to the range of colors we can see, the details we can visually resolve, the sounds we can hear, and the odors we can distinguish, limits that are also set by our physiology. We don't have difficulty accepting these limitations either. Why then should we be less willing to accept the existence of limits on the way we think about things? In fact, the most exciting thing about studying contemporary physics is learning how we have come to accept our limitations and arrive at an understanding of nature in spite of them.

Morton Tavel, *Contemporary Physics and the Limits of Knowledge* (Rutgers, 2002), p. 16.