

Matter, Mind and Mechanics

New models for dynamogenesis and rationality

Synopsis of *The Mechanical Foundations of Psychology and Economics*

Jon Doyle

Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
doyle@mit.edu

October 2, 2000

Revised draft of February 13, 2001 at 15:47

The mathematical concepts of modern axiomatic rational mechanics apply more broadly than generally recognized, in that formal mechanical concepts apply directly and with slight adaptations to certain psychological and economic systems as well the familiar physical applications. These nonphysical applications provide new means for characterizing realistic notions of economic rationality and limits on reasoning abilities, translate psychology and economics into studies of new types of mechanical materials, and open traditional informal philosophical conceptions of materialism and dualism to new avenues of mathematical and experimental investigation. This semi-formal article, which summarizes a detailed book-length formal treatment [13], briefly sketches modern mechanics, the adaptation of mechanical axioms to cover hybrid and discrete systems, an illustrative formalization of a representative rational psychological system from artificial intelligence, and the new perspective on psychology and economics, along with reflections on related issues in mechanical computation and biology.

Copyright © 2000, 2001 by Jon Doyle. All rights reserved.

Contents

1	Introduction	4
1.1	Rethinking materialism	5
1.2	Characterizing rationality	6
1.3	Understanding character	7
1.4	Method and structure of the presentation	8
2	The structure of modern rational mechanics	9
2.1	General laws	10
2.2	Special laws	12
3	Hybrid mechanical systems	13
3.1	Structures of hybrid mechanics	14
3.2	Illustration: mind and body	16
4	Discrete mechanical systems	18
4.1	Structures of discrete mechanics	19
4.2	Illustration: reason maintenance	22
4.2.1	Psychological structure and behavior	22
4.2.2	Mechanical content and formalization	24
5	Mechanical rationality	26
5.1	Structure of economic rationality	27
5.2	Limitations on rationality	29
5.2.1	Kinematic limitations	30
5.2.2	Dynamic limitations	32
6	Discussion	38
6.1	Defining mechanics	38
6.2	Prediction and design	39
6.3	Open problems	40
6.4	Precursors	41
7	Conclusion	44
	Acknowledgment	45
	References	45

In memoriam

Clifford Ambrose Truesdell III
February 18, 1919—January 14, 2000

Herbert Alexander Simon
June 15, 1916—February 9, 2001

1 Introduction

The advance of science accelerated dramatically in the seventeenth and eighteenth centuries when discoveries in mathematics and rational mechanics by Newton, Euler, and others changed the character of natural philosophy from informal, largely philosophical debate to technical investigation and explicit calculation of the properties of physical systems modeled as evolving in accordance with specific mathematical equations and starting from specific hypothesized initial conditions. These mathematical models enabled scientists to refine physical theories, and enabled engineers to construct complicated physical systems to meet precise specifications. The advance in understanding changed perspectives so much that Leibniz claimed sufficiently great calculating abilities and a full description of initial conditions would permit determination of the entire future of the world subsequent to the initial time.

Such optimism did not extend to the contemporaneous early stages of the human sciences of psychology and economics. The seventeenth century also saw Descartes' promulgation of a dualistic theory of mind, in which a mental substance of the mind accompanied the physical substance of the body. Discourse at the time also spoke of forces on minds and bodies, just as it does today. In spite of such conceptions of mind and body consisting of substances acted upon by forces, the mathematical tools of the new mechanics did not apply to Cartesian minds, for their mental substances lacked physical position, meaning that mental actions lacked description in terms of the physical motion treated by mechanics. The new mechanics thus offered no way to apply its developing formal concepts to understanding the relation of minds to the body or the nature of forces acting on minds, leaving mathematical theories of psychology and economics to emerge later from non-mechanical theories of logic, probability, and utility. This difference with respect to mechanical formalism produced an increasingly wide separation of the mental and the physical sciences.

The scientific import of this divorce of mental and physical sciences became clearer later as psychology began to explore computational characterizations of reasoning and behavior, and as economics began to cast about for theories better matching human capabilities than its foundational theory of ideal rational choice. Computational formalizations of psychological theories involved motion in spaces of mental states which, though very different than physical space, at least proved susceptible of mathematical formalization. Realistic economists grew appreciative of the hard work involved in making choices and of the slowness of the mind to change when subjected to new information or other influences. Popular discourse still spoke of mental forces, work, and inertia to reflect these concerns, much as in the days of Descartes and Newton. People also came to use mechanical concepts of inertia, force, energy, and pressure informally in describing economic markets and behavior. In spite of the continuing application of seemingly similar concepts, the divorce of the mental and physical sciences impoverished the mental sciences when compared to the physical sciences by abandoning to the purely physical realm mechanical concepts of force and inertia that proved fruitful in analyzing physical behaviors. Study a physical problem, and one has recourse to physics, chemistry, and biology, as well as differential equations and mathematical theorems that aid in analysis and prediction. Study a mental system, and one lacks almost all of this intellectual heritage, for the traditional conceptual tools do not apply.

The scientific separation of mental and physical need not stand. We bridge the gap between matter and mind with mechanics. We indicate how taking a closer look at the quiet progress of mechanics in recent

years provides formal concepts of force, mass, momentum and work that enable one to transform some heretofore metaphorical uses of these terms into meaningful, true or false, non-metaphorical statements within the axiomatic framework of modern rational mechanics. This augmentation of the existing technical conceptions of logic, economics, and computational intelligence with the formal concepts of mechanics permits construction of mechanical theories of the interaction of mind and body, limits on ideal economic rationality, and the character of everyday psychology.

The resulting reconciliation of the mental and the physical opens some long-standing philosophical problems to serious mathematical investigation and so improves prospects for removing the technical limitations handicapping the human sciences relative to the physical sciences.

1.1 Rethinking materialism

Philosophers have speculated for centuries about possible relations between mind and body, with theories ranging from nonexistence of mind to nonexistence of body, and from complete disconnection of mind from body to complete correspondence of mind and body.

Although Decartes and others viewed mind and body as somewhat separate entities acting on each other, such theories fell into disrepute for at least two reasons. First, proponents of such dualistic theories could not supply any formal model for or rules governing the proposed interactions. Science was just beginning to understand physical forces in mathematical terms, but not in a way that applied to understanding interactions of mind and body. Second, even setting aside the lack of a formal model for mind-body forces, proponents could not identify measurable forces of this type. Everyone knew that lifting an arm caused the arm to rise, but this apparently causal relation did not seem to fit the mechanical mold because the mental effort required to lift the arm did not vary with the weight of things held by the arm.

The demise of dualism and the everyday incredibility of idealism have facilitated the dominance of the scientific viewpoint by materialism, which today we interpret as the view that everything in the world consists of the material particles and fields of physics. Materialism even dominates thinking in psychology, where a tradition of behaviorism, computational mechanism, and neurophysiological primacy sidelines the rather obvious disconnection between the experience of mental life and the specific circumstances of embodiment. Though many have observed the tenuous and inessential grounding of commonsense psychological concepts in the physical, chemical, and neurophysiological bodily base, better understanding of the dramatic effects of certain chemicals and bodily trauma on mental behavior have lead others to dismiss ordinary psychological talk about beliefs, desires, emotions, volition and the like as hopelessly flawed, even as an approximation to some supposed neurophysiological truth. This philosophical attitude has developed along with a shift in emphasis in artificial intelligence away from symbolic models of thought toward numerical models that overlay varying degrees of non-numerical structure.

Economics, in contrast, represents perhaps the last holdout of amaterialist thinking, even though it epitomizes “materialistic thinking” in an irrelevant popular sense. Theoretical economics deals mainly with a theory of ideal rationality based solely on belief, preference, and choice independent of any materialist grounding. The preferences economists attribute to humans may derive from complex balances among underlying motivations, emotions, and sentiments, but these underlying origins do not concern the economist

as long as these origins yield preferences fitting the form demanded by ideal rationality.

The present work provides both formal means to rehabilitate dualistic psychologies and motivations for doing so. We use the modern axioms of mechanics to characterize the structure of forces exerted by mind on body and by body on mind in exactly the same way that these axioms characterize forces among physical bodies. Though the early formulations of mechanics did not cover the Cartesian conception of mental substances, the mathematically refined axioms developed in the twentieth century provide the necessary breadth. The resulting mechanical characterization of mental substances calls into question the conventional conception of materialism, much as discoveries in physics have extended the ancient concept of materialism to include the invisible fields of electromagnetism and various quantum particle fields unobservable in everyday life, and to include a theory of gravitation in which energy itself has mass.

By providing a rigorous conception of forces between mind and body, the present work provides new avenues for exploring questions about the existence, nature, and origin of such forces. The formal framework covers many possibilities with no presumption that any or all of these exist in our world. It provides a formalization of interaction in which the mind raises the arm by forcing communication of information that initiates a physical action. The formal framework provided here also makes sense of telekinesis in the same way that traditional mechanics makes room for all sorts of possible forces not known to the ancients. We make no suppositions regarding telekinesis or other undemonstrated forces and leave questions about the existence additional forces beyond those commonly recognized to future physical and psychological investigation.

1.2 Characterizing rationality

The economic ideal of rational choice enjoys a deep mathematical development unrivaled by any science outside the physical sciences save for logic and the theory of computation. This very theoretical strength also reflects the greatest weakness of economics, namely the mismatch between the ideal of economic rationality and the reality of decision making by humans of limited mental capacity and rationality. Although economics posits beliefs and preferences that exhibit unbounded degrees of accuracy, completeness, and consistency at each instant of acting, humans notoriously find thinking difficult. They have trouble seeing the consequences of their beliefs and preferences for the decisions that economics supposes they rationally make. Even when they come to rational decisions, impulsive desires can interfere and cause them to act against their rational judgments. Confronted with new information, they often fail to adjust their expectations in the comprehensive way prescribed by economics, or can take a long time to do so.

Recognizing the mismatch between theory and actuality, thoughtful economists have sought theories of more realistic conceptions rationality for some time. This search has met with only modest success because it implies development of a respectable mathematical model of human or human-like thinking. Economics should find such a model in psychology, but psychology has no such model to provide. Neurophysiological models exist in mathematical form, but the mathematical concepts they involve say nothing about the properties of interest to the economist. The non-neuronal symbolic and structured probabilistic models of artificial intelligence provide detailed, concrete, and formal models much closer to the needs of economics, but these models of specific types of thinking do not yet include the mathematical models of limits to rea-

soning needed by the economists.

Numerous truisms attest to the limits of economic formalism and point to a solution. Everybody knows from self-help books that you have to force yourself to change, and the bigger the desired change, the more you have to work at it. Everybody knows that the more you know about some question the harder it is to change your position. Everybody knows that getting going on tackling a task is often the hardest problem faced in carrying out the task, and that once one overcomes this initial inertia progress generates its own momentum that keeps one going. Everybody knows that maintaining direction or focus requires forcing oneself to ignore distractions. One can speculate that people express and use such truisms because people have well-developed abilities for predicting the motions of physical objects in such terms and because everyone gains some understanding by applying the same terms and abilities to mental behavior, even after acknowledging the unpredictable ways minds can act compared to inanimate objects.

Science, when confronted by things “everybody knows”, can seek to confirm, disprove, refine, or correct the claims experimentally or theoretically. Doing any of these requires making the claims precise, preferably in measurable ways. If scientists judge they can never make the claims precise, they may dismiss them as nonsense or misunderstandings. Scientists have regarded many psychological truisms as nonsense for centuries because the mathematical formalism developed in physics have not applied in clear ways to mental “forces” or “inertia”, even though truisms about these make sense in everyday life and seem consonant with physical intuition. This creates an uncomfortable conflict for the humble and reflective scientist, since everyday living involves applying terms one disbelieves intellectually.

The mechanics sketched here points a way to reducing these conflicts. We certainly do not yet know how to make all such truisms and beliefs sensible, and some might remain formally uninterpretable forever, but the mechanics described here indicates how to transform some truisms and common beliefs about psychology into sensible formal statements of mechanics which one might show true or false. Some truisms correspond to true formal statements, justifying their persistence despite centuries of scientific approbation. These formalizations identify notions of psychological and economic force, mass, and momentum as precise as the familiar physical ones and obeying essentially the same laws and interrelations.

Identifying mechanical forces and masses in psychology and economics provides a new formal vocabulary for characterizing limits on rationality. This new vocabulary uses mass and force to express limits on the speed with which agents can change mental state and direction in accommodating new information and in reasoning and deliberation. It analyzes difficulties in maintaining focus of attention in terms of the magnitude and direction of forces resulting from superposition of competing motivations, and in terms of the work required to counteract distracting forces. It views some equilibrium notions of economics in terms of static balance of forces and relaxed or equilibrium states of materials. Mechanical limitations on the rapidity and effort of change and the maintainance of attention do not account for all limitations on rationality, but they represent some of the most important limitations that current psychological theories characterize poorly.

1.3 Understanding character

Rationality constitutes just one of the notable characteristics of humans. As noted earlier, the economic conception of rationality abstracts from the many sources of belief and preference in the underlying psychology

of the agent. Some of these underlying sources vary from agent to agent and from time to time, but others characterize different personalities or psychological types. The mechanical perspective offers additional insight in understanding such classifications of agents through the notion of distinct mechanical materials.

The underlying native mental competences and inherent motivations and proclivities that distinguish people do not correspond to any known properties of neurons or even assemblages of neurons, but do correspond in fairly direct ways to material characterizations familiar in mechanics. Some important components of native competence find reflection in kinematic assumptions that restrict mental states to ones exhibiting certain degrees of local consistency and completeness in the same way that rigidity assumptions restrict material states to ones exhibiting certain relationships among distances between portions of bodies. Similarly, important characterizations of inherent motivations and reasoning processes, such as common human drives and variations in the deliberativeness, wantonness, and conservatism exhibited in different personality types, in turn find reflection in assumptions about the forces generated by and acting in different types of persons and their behaviors. The mechanical interpretation sketched here treats personality types and more refined classifications of human character on a par with mechanical identifications of elastic, electromagnetic, and gravitational materials.

1.4 Method and structure of the presentation

We describe mechanical models for psychology and economics in several steps. The descriptions presented in here merely summarize the detailed mathematical constructions and analyses of examples presented in the full book [13].

First, we summarize the structure and content of the axioms of modern rational mechanics. Most physics textbooks do not present axioms, or present only Newton's informal and conceptually-inadequate postulates. The modern axioms reveal the true generality of mechanics in ways that usual textbook presentations do not.

Second, we use these axiomatic concepts to define a notion of hybrid mechanics, in which two or more mechanical worlds develop concurrently and interact (or not) through forces. Hybrid mechanics follows standard mechanics in almost every respect. It provides a straightforward, coherent model for interaction of matter and mind through forces obeying mechanical laws.

Third, we broaden the axioms for individual mechanical systems by removing requirements that space and time form continua. This broadening preserves almost all the familiar structure of mechanics in a way that permits direct mechanical formalization of standard discrete or symbolic notions of psychology. We then illustrate discrete hybrid mechanical characterizations of psychological and economic systems by sketching a mechanical formalization of an artificial psychology from artificial intelligence.

Fourth, we summarize the formal structure of the economic theory of ideal rationality, and sketch the use of mechanical concepts in formalizing new conceptions of limited rationality.

We conclude by discussing some methodological, technical, mathematical, and philosophical problems opened up by these models, and by comparing characteristics of the present formalization against the reasons underlying the failures of past attempts at understanding the mind in physical terms.

2 The structure of modern rational mechanics

[T]he aim of theoretical physics is to construct mathematical models such as to enable us, from the use of knowledge gathered in a few observations, to predict by logical processes the outcomes in many other circumstances. Any logically sound theory satisfying this condition is a good theory, whether or not it be derived from “ultimate” or “fundamental” truth. It is as ridiculous to deride continuum physics because it is not obtained from nuclear physics as it would be to reproach it with lack of foundation in the Bible.

C. Truesdell and W. Noll, [59, pp. 2–3]

We broaden mechanics by starting with the axioms formulated in modern rational mechanics, especially those developed by Noll [35, 36, 38, 37, 41] and expounded by Truesdell [60]. Through clear formulation in terms of modern mathematical concepts, these axioms already distinguish many central mechanical concepts from special characteristics of physical space and time, and so reduce the revisions needed to adapt the axioms to cover discrete psychological systems. A contemporaneous logical axiomatization due to Suppes and his collaborators [32] offers a less attractive alternative for this purpose because it hews closely to the numerical formulations that dominate elementary textbooks. One need not take a stand on Truesdell’s [58] criticisms of that axiomatization to appreciate that formulating mechanical notions directly in terms of continuous number systems makes adaptation of mechanics to psychology much harder by hiding mechanical essence within a sheathing of numerical accident.

Although the axioms on which we build have enjoyed widespread use for decades among mathematicians and mechanical engineers, they differ from what most people think of as “the” axioms of mechanics. Many have heard that Newton propounded three axioms, and might even remember translations or paraphrases of their statements from the Latin. Readers of introductory physics texts might recall the equation $f = ma$ as the formal embodiment of mechanics. Neither of these candidates deserves coronation as the axioms of mechanics, however. Newton’s postulates lack any precise formal meaning, and have little to do with what modern physics reads into them. By itself, $f = ma$ provides no axiomatic basis for mechanics, though it appears as an expression within one axiom among many in some modern formulations. These candidates fail as axioms in many ways, but most importantly in providing no characterization of the notion of force. Adequate axioms for forces were developed first in the 1950s.

Modern theories of mechanics regard bodies as subject to general laws applying to all types of materials, laws that relate forces on bodies to the motions caused by the forces. None of the most general laws say anything about which forces exist, or even that *any* particular forces exist. Such statements instead come in special laws of *dynamogenesis* that characterize the origin of forces. Other special laws characterize the behavior of special types of materials, ordinarily identified in terms of constraints on bodies, configurations, motions, and forces. For example, rigid body mechanics comes from adding to the general theory kinematical constraints that fix the relative distances of body parts, while the theory of rubber comes from augmenting the general theory with the configuration-dependent forces characteristic of rubber. Mechanics uses the term *constitutive assumption* to refer to the special laws for particular materials, since the laws reflect assumptions about the constitution of the material. Mechanical practice depends critically on these

special laws. So-called fundamental laws of physics stand largely irrelevant to the special laws, as rigorous derivation of most special laws from “fundamental” properties of elementary particles remains well beyond present theoretical capabilities.

2.1 General laws

Noll’s general axioms break down into axioms characterizing geometry, bodies, and forces. We summarize the axioms and constructions regarding these topics in turn, mainly following Noll’s exposition in [38].

Time and space: Noll’s first set of axioms characterizes *neo-classical event worlds* of time, space, duration and distance. One can vary this development to obtain axioms for relativistic event worlds, but we avoid that here as a complication irrelevant to explaining the central concepts of hybrid mechanics.

We start with a set \mathcal{W} of *events* called the *event world*. We first add to the event world a *time-lapse* function $\hat{t} : \mathcal{W} \times \mathcal{W} \rightarrow \mathbb{R}$ giving the length or duration of the temporal interval between events, where $\hat{t}(e, e') > 0$ means e occurs before e' . Three axioms characterize this function as antisymmetric under time reversal, additive for intervals sharing a common event endpoint, and reflecting a continuum of events and times.

The time-lapse function induces an equivalence relation $\Sigma_{\mathcal{W}, \hat{t}} \subseteq \mathcal{W} \times \mathcal{W}$ (abbreviated Σ) of *simultaneity* on \mathcal{W} and thus a partition $\Gamma_{\mathcal{W}, \hat{t}} \subset 2^{\mathcal{W}}$ (abbreviated Γ) of \mathcal{W} into sets of simultaneous events called *instants*. From \hat{t} and Γ one obtains a time-lapse function $\bar{t} : \Gamma \times \Gamma \rightarrow \mathbb{R}$ over instants. The absolute value of \bar{t} provides a Euclidean metric on Γ .

We next add to the event world a *distance* function $\hat{d} : \Sigma \rightarrow \mathbb{R}^{0+}$ taking pairs of simultaneous events to nonnegative numbers. Two axioms characterize this distance function as inducing a Euclidean metric on the events in each instant $\tau \in \Gamma$ that makes the translation space of each instant into a three-dimensional vector space isomorphic to a vector space \mathcal{V} itself isomorphic to \mathbb{R}^3 . Together, these axioms on time and distance characterize neo-classical event worlds.

Bodies: Noll’s second set of axioms characterizes *material universes* of bodies and their parts. A material universe consists of a set Υ of *bodies* together with a *subbody* or part-of relation \preceq on $\Upsilon \times \Upsilon$. We read $\mathcal{B} \preceq \mathcal{C}$ as saying that \mathcal{B} is a (possibly improper) subbody or part of \mathcal{C} . Six axioms stating antisymmetry and transitivity of the subbody relation, the existence of greatest and least bodies, the existence of a unique exterior \mathcal{B}^e of each body $\mathcal{B} \in \Upsilon$, the separateness of body parts from the body exterior, and the existence of meets characterize the material universe, and together imply that the material universe forms a boolean lattice.

We follow almost all treatments of mechanics and assume for simplicity that each body consists of a set of *body points*. Bodies consisting of single points can represent the point-masses of analytical mechanics, while bodies consisting of continua of points of different dimensions represent the solids, shells, rods, and other bodies studied in continuum mechanics. The assumption that bodies consist of sets implies the material universe forms a field of sets.

Motions: Noll uses the preceding structures for bodies and events to construct straightforward kinematical notions. Noll takes care to separate the reality under analysis from mathematical descriptions of that reality in terms of coordinate systems. In particular, he analyzes behavior in terms of configurations and deformation processes. He defines *configurations* of a body intrinsically in terms of the sets of distances between points of the body and *deformation processes* as histories of configurations over temporal intervals.

With behavior described in realistic (coordinate-free) terms, Noll proceeds to describe motions with respect to particular coordinate systems or frames of reference in terms of placements and time-indexed families of placements. *Placements* of bodies in space consist of mappings of body points into points of space. Standard mechanics assumes smoothness of these placements. *Motions* of bodies consist of a family of placements into spacetime instants. *Frames of reference* or reference processes consist of families of Euclidean metrics on instants. Noll defines rates of motion and motion and rates relative to frames of reference in straightforward ways. Standard mechanics also assumes continuity of the first two derivatives of motions and frames of reference.

Forces: Noll's third set of axioms characterizes the forces existing at each instant in terms of *force systems*. A force system consists of a mapping $f : \Upsilon \times \Upsilon \rightarrow \mathcal{V}$ of pairs of bodies to spatial vectors, where $f(\mathcal{B}, \mathcal{C})$ denotes the force exerted on \mathcal{B} by \mathcal{C} in the system of forces f . A system of *torques* consists of a similar mapping of separate bodies and spatial locations to linear transformations of spatial vectors.

Noll's writings state mechanical axioms about forces in several equivalent ways. We here describe a presentation of the axioms by Truesdell [60] that separates the central structures from the specific assumptions of continuum mechanics somewhat more explicitly than do some of Noll's continuum-focussed developments.

The first axiom of forces states that forces exerted on or by separate subbodies of a body combine additively. The second axiom states a balance of forces, expressed through the requirement that the exterior of a body exerts no force on the body, that the sum of all forces on each body vanishes. This last condition may seem strange, but really just amounts to treating all components of familiar force equations equally. In particular, this axiom leads to treating the inertial force $-ma$ as just one special sort of force cancelling out the sum of all the other forces on the body, thus rewriting the familiar $f = ma$ as $f - ma = 0$. The third axiom states that a similar balance of torques holds for all bodies and all spatial reference points. These laws apply more generally than might appear, since they hold for all bodies and parts of bodies, not merely body-points.

Each force system induces two subsidiary force systems, internal forces between separate parts of a body, and external forces exerted by the exterior of a body on its parts. The fourth axiom states that the internal forces are contact forces that vary continuously with the area of the contact boundary between the two parts, while the fifth axiom states that the external forces can be by contact and at a distance and vary continuously with the mass, volume and contact boundary areas of the parts. Standard continuum mechanics shows that these axioms on body and contact forces imply the existence of Cauchy's stress tensor, which in turn summarizes both types of forces on the body through Cauchy's equation of motion $\dot{p} = B + \text{div}(T)$, in which $p = mv$ is the inertia, $-\dot{p}$ is the inertial force, B is the body force, and $\text{div}(T)$ is the divergence of

the stress tensor T .

Dynamogenesis: The preceding axioms characterize the structure of systems of forces in isolation, but say nothing about how forces arise in the evolution of mechanical systems. Although most specific characterizations of forces depend on the specific class of material involved, Noll states three additional axioms concerning dynamogenesis. The first of these axioms states the principle of *determinism*, that the history of body and contact forces (or equivalently, the stress) at preceding instants determines a unique value for these forces at a given instant. The second axiom states the principle of *locality*, that the forces at a given point depend only on the configuration of bodies within arbitrarily small neighborhoods of the point. The third axiom states the principle of *frame indifference*, that forces depend only on the intrinsic properties of motions and deformation, not on properties that vary with the reference frame. Frame indifference thus distinguishes the underlying reality (motion) from particular ways of describing that motion (coordinates and speeds in reference frames).

Energy: Noll postulates a measure $e : \Upsilon \rightarrow \mathbb{R}^{0+}$ assigning to each body its nonnegative *internal energy*, together with a mapping $q : \Upsilon \times \Upsilon \rightarrow \mathbb{R}$ called a system of *heatings*. The system of heatings satisfies axioms similar to those defining systems of forces and others that require that heatings correspond to transfers of energy between bodies.

The mathematical fiction of potential energy, so useful in analyzing motion in gravitational and electromagnetic fields, plays no axiomatic role in mechanics and bears no relation to the mechanical notion of internal energy.

Noll defines *work* in terms of force and distance as usual, and proves the balance of forces and torques corresponds exactly to frame indifference of the rate of work done by bodies.

2.2 Special laws

In themselves, the general laws provide a weak basis for mechanical prediction because they make no assumptions about the properties of specific materials. The builder's choice of reinforcing steel over bamboo, the auto mechanic's choice of lubricating motor oil over sand, and the magician's choice of levitating magnetic steel over wooden containers all exploit the different characteristics of different materials. Mechanics thus incorporates many different specialized axioms characterizing the properties of special materials. These constitutive assumptions may characterize broad classes and abstract properties, or narrow and specific ones.

Geometry: Some important constitutive assumptions characterize materials in terms of their geometric structure. Characterizations of rigid bodies, solids, fluids, crystals, liquid crystals, incompressible materials, and isotropic materials all involve kinematic assumptions about the possible material geometry and motion, either as the entire assumptions or as key elements of such assumptions.

Dynamogenesis: Some important constitutive assumptions characterize materials in terms of specific force laws or of characteristics of the force laws. General examples here include elastic materials and materials with fading memory. More specific examples include materials generating gravitational, electromagnetic, or inertial forces.

Mass: Inertial forces represent the most basic of the most general constitutive assumptions concerning fundamental physical forces. As usual, Noll postulates a measure $m : \Upsilon \rightarrow \mathbb{R}^{0+}$ assigning to each body its nonnegative inertial mass value. He defines momentum in a reference process as the sum or integral of velocity with respect to mass. In the case of point bodies, this identifies the momentum as the usual product of the mass and velocity. Noll notes the usual theoretical distinction between gravitational and inertial mass but drops it in the formalism for the usual reason of empirical identity.

Inertial forces: Truesdell [60] formulates the mechanical laws of inertial forces in two axioms. The first axiom states the existence of inertial frames, frames of reference in which a body has constant momentum during some interval if and only if the resultant force on the body vanishes. The second axiom refines the first to say that in an inertial frame the resultant force on a body equals the negative derivative of the momentum of the body. Combining these axioms with the general balance laws of forces and torques yields Euler's fundamental laws $f = \dot{p}$ and $F = \dot{M}$ stating the respective balances of linear and rotational momentum.

Noll [40] regards the axioms of inertia as special laws in part because forces generated by masses play an important role in many, but not all, mechanical problems. He cites the motion of toothpaste out of a squeezed tube as an example showing the occasional effective irrelevance of mass to motion.

Viewing the laws of inertia as special rather than general laws bears even more significance than mere occasional irrelevance of mass would suggest. Observing that motion may occur independent of mass provides a precedent for physical mass that need not always play a role in mental motion, and for mental mass that need not always play a role in physical motion.

3 Hybrid mechanical systems

Recent work on so-called hybrid system models has provided mathematical tools for describing systems exhibiting both discrete and continuous behaviors [2, 7, 9]. The simplest such models characterize the system motion in terms of the product of a finite automaton and a space of smooth motions. Each state of the automaton corresponds to a different smooth dynamics with its own acceptable operating region in the continuous portion of the state space. Smooth motions that move outside the acceptable region trigger transitions between the discrete states. The nature of the triggering conditions and the automaton structure guide the actual discrete state transitions. The smooth motion continues under the regime imposed by the new discrete state. For example, the pendulum clock corresponds to a fairly trivial system in which the smooth dynamics remain the same in all discrete states, but in which movement of the pendulum to certain angles triggers transitions from one gear state to the next, with groups of these gear states representing the times displayed by the clock through its second, minute, and hour hand positions.

These hybrid models, though useful in providing conceptual tools for analyzing complex systems, do

not in themselves restrict system structure or behavior in any way, since the hybrid models themselves may involve any possible discrete and continuous structures. In particular, hybrid models by themselves provide no means for reconciling psychology and economics with physical law. To take the first step toward a reconciliation, we extend the axiomatic concepts of traditional mechanics to define the notion of hybrid mechanics. The notions of geometry and force provide the foundations of hybrid mechanics, just as in the mechanics of Newton and Noll. Roughly speaking, a hybrid mechanical system exhibits a spatial geometry obtained as the product of the spatial geometries of the factor mechanical systems, and a space of forces obtained as the sum of the force-spaces of the factor systems. The factor systems interact through the forces characteristic of the hybrid materials, or more specifically, through the way the the product states give rise to the product forces. This formulation permits, but does not require, each factor system to contribute to the forces acting in other factor systems by the hybrid dynamogenesis relation.

In this section, we sketch the notion of hybrid mechanics for the case in which one combines several traditional mechanical systems into one hybrid system. This case of hybrid mechanics stays close to traditional mechanical concepts. Indeed, we might view one of Noll's constructions, in which he takes a full Euclidean vector space to represent the exterior of a body [38, p. 77], as a very simple instance of a hybrid mechanics. One may similarly interpret a brief query Truesdell once posed seeking evidence of measurable nonphysical forces on material objects [56], as indicating he considered mechanics as extending to cover potential interactions between the physical and the mental (though in fairness, Truesdell usually mentioned psychology only to disparage the intellectual shoddiness of that field in comparison with the rigor attained in mechanics and other sciences).

3.1 Structures of hybrid mechanics

We start with a set $\Xi = \{\xi_i \mid i \in \mathcal{I}\}$ of mechanical systems indexed by some index set \mathcal{I} . We characterize hybrid time, space, bodies, forces, and kinematical and dynamical motions in terms of the component mechanical structures.

Time and space: We characterize a hybrid event world over Ξ in terms of a set \mathcal{W} , functions \hat{t} and \bar{t} , and vector space \mathcal{V} related to the corresponding concepts of the component worlds. Define $\vec{\mathcal{W}} = \prod_i \mathcal{W}_i$ as the product of these sets of events and $\check{\mathcal{W}} = \bigoplus_i \mathcal{W}_i$ as their disjoint union. We first require that hybrid spacetime events consist of events from each of the component event worlds, or stated formally, $\mathcal{W} \subseteq \vec{\mathcal{W}}$. Not all combinations of component events makes up a hybrid event, only those that make sense given the time-lapse functions on the component spacetimes and the way we match up simultaneous events in different component spacetimes. Accordingly we require that the hybrid time lapse function agree in sense and simultaneity with the component time lapse functions. We insist that all component spacetime events fit in somewhere in the hybrid spacetime by requiring for each $e \in \check{\mathcal{W}}$ there exist $\vec{e} \in \vec{\mathcal{W}}$ and $i \in \mathcal{I}$ such that $e = \vec{e}_i$. As with the component spacetimes, the time-lapse function \hat{t} induces a simultaneity relation Σ on \mathcal{W} , which in turn induces a partition Γ of \mathcal{W} . Our next restriction on \mathcal{W} requires that hybrid instants look like products of component instants. Let $\vec{\Gamma} = \prod_i \Gamma_i$. We then require that for each $\tau \in \Gamma$ there exists some $\vec{\tau} \in \vec{\Gamma}$ such that $\tau = \prod_i \vec{\tau}_i$. Since each component instant τ_i is isomorphic to a vector space \mathcal{V}_i , this means

that each hybrid instant is isomorphic to the product vector space $\mathcal{V} = \prod_i \mathcal{V}_i$. This combination of the component time metrics permits the hybrid mechanics to scale time differently than any of the component time metrics.

We make no immediate use of the hybrid distance metric \hat{d} on \mathcal{V} , and require only its compatibility with the component distance metrics.

Bodies: We obtain the bodies of the hybrid mechanics by simple combination of the bodies of the component mechanics. Each component universe of bodies Υ_i comprises a boolean lattice, and we form the hybrid universe $\Upsilon = \bigoplus_i \Upsilon_i$ by taking the direct sum of these lattices. This direct sum also forms a boolean lattice. For boolean lattices the direct sum and direct product have the same result, with the hybrid body \mathcal{B}_i for $\mathcal{B}_i \in \Upsilon_i$ represented in the product by an element \mathcal{B}' such that $\mathcal{B}'_i = \mathcal{B}_i$ and $\mathcal{B}'_j = \top_{\Upsilon_j}$ for each $j \in \mathcal{I} \setminus \{i\}$, where \top_{Υ_j} stands for the maximum element in Υ_j . We use the sum notation in preference to the product notation because it usually makes it easier to discuss forces among bodies of different type.

Our assumption that each component universe consists of a boolean lattice of sets of point bodies permits us to view each body in the hybrid universe as a set of point bodies as well.

Motions: The kinematic definition of motion in a hybrid mechanics follows that of motion in the component mechanics. A motion consists of a mapping $\chi : \mathcal{B} \times I \rightarrow 2^{\mathcal{W}}$ from body points in a body \mathcal{B} and instants in some temporal interval I to sets of events such that $\chi(\mathcal{C}, \tau) \in \tau$ for each point body $\mathcal{C} \in \mathcal{B}$ and instant $\tau \in I$, that is, a mapping of body points to events giving the placement of the body point in each instant. We write χ_i to denote the projection of a hybrid motion χ onto the event world \mathcal{W}_i .

Mass: Hybrid mechanics generalizes the traditional conception of mass directly by regarding hybrid mass as the vector of mass values in each of the component systems. Newtonian mechanics already recognized conceptually different types of mass in inertial mass and gravitational mass. The only theoretical difficulty posed by working with these two different types of mass was not mathematical but instead explaining the experimental equivalence of these concepts in all known practical circumstances, a question answered by Einstein in his theory of general relativity. The vectorial conception of mass used in hybrid mechanics represents theoretically similar differences in categories of mass.

Forces: Force systems in hybrid mechanics take exactly the same form as the force systems in the component mechanics, namely a mapping $f : \Upsilon \times \Upsilon \rightarrow \mathcal{V}$ of pairs of bodies to spatial vectors. As in Noll's axioms, we require the hybrid force system to exhibit additivity on separate bodies, and in particular, on bodies from different components. We require the same balance of forces as in Noll's axioms, and the same continuity and boundedness constraints on body and contact forces within each factor mechanics. Indeed, regarding mass as a vector concept means the familiar equation $f = ma$ holds in each component system when we restrict attention to the projections f_i , m_i , and a_i of the values of a hybrid force system, masses, and accelerations into the values holding within each component mechanics. The differences between hybrid mechanics and traditional mechanics lie elsewhere.

The first difference comes in considering the balance of torques posited by Noll. The notion of torque makes sense in each of the component spacetimes, but in the hybrid spacetime we must interpret torques as vectors of component torques, not as general linear transformations over the hybrid space. To retain the balance of torques, we insist that torques balance within each component system given the force value projections f_i .

Dynamogenesis: The larger differences between hybrid and traditional mechanics appear in considering Noll's axioms concerning dynamogenesis. His axiom of frame indifference carries over unchanged in the hybrid context, as one might expect, but the axioms of determinism and locality require reformulation.

In hybrid mechanics, the force on a body within some component system potentially depends not only on contributions exerted by other bodies within the same component system, but also on bodies in other component systems. This means that axioms of determinism and locality cannot hold within the component systems, unless the separate mechanical components are to evolve separately with no interaction. Accordingly, we state the general axioms of determinism and locality at the hybrid level, as restrictions on the systems of hybrid forces, not as restrictions on the component force systems. Naturally, specific types of hybrid systems might reintroduce determinism and locality constraints on the component mechanics to capture the characteristics of special materials.

3.2 Illustration: mind and body

We present here a simple illustration of how to use the preceding formal framework to address in a straightforward technical manner psychological theories of the interaction and correlation of mind and body, assuming temporarily, for the purpose of illustration, that the mind constitutes just another standard physical system.

For this illustration, we employ a hybrid mechanics with two component worlds, one physical, one mental, to model the body and mind of a person we will call René. For the physical world we take the standard mechanical model, and identify the body of René as a body \mathcal{B}_p existing within a universe Υ_p of physical bodies. Similarly, for the mental world we identify the mind of René as a body \mathcal{B}_m existing within a separate universe Υ_m of mental bodies. Our simple illustration then uses the hybrid body $\mathcal{B}_R = \mathcal{B}_p + \mathcal{B}_m = \mathcal{B}_p \sqcup \mathcal{B}_m$ to model the person René, and the hybrid body $\top = \top_p + \top_m = \top_p \sqcup \top_m$ to model the universal body.

Interaction: Let us now consider the total force acting on René at some instant. In the hybrid mechanical formalism, we write this force as $f(\mathcal{B}_R, \top)$. The separation of the mental and physical bodies and the axioms for forces then let us rewrite this force as

$$f(\mathcal{B}_R, \top) = f(\mathcal{B}_p, \top) + f(\mathcal{B}_m, \top). \quad (1)$$

Let us look first at $f(\mathcal{B}_p, \top)$, which represents the force on the physical body. We decompose \top into

separate components

$$\top = \mathcal{B}_R + \mathcal{B}_R^e = \mathcal{B}_p + \mathcal{B}_m + \mathcal{B}_p^{pe} + \mathcal{B}_m^{me}, \quad (2)$$

where \mathcal{B}_p^{pe} denotes the physical environment of the physical body obtained as the relative complement of \mathcal{B}_p with respect to the greatest physical body \top_p , and where \mathcal{B}_m^{me} denotes the mental environment of the mind obtained as the relative complement of \mathcal{B}_m with respect to the greatest mental body \top_m . With this partition of the hybrid universe, we apply the axioms for forces to rewrite the force on the physical body as

$$f(\mathcal{B}_p, \top) = f(\mathcal{B}_p, \mathcal{B}_p) + f(\mathcal{B}_p, \mathcal{B}_m) + f(\mathcal{B}_p, \mathcal{B}_p^{pe}) + f(\mathcal{B}_p, \mathcal{B}_m^{me}). \quad (3)$$

This just says the total force on the physical body consists of the sum of the forces exerted on the body by the body itself, by the mind, by the physical environment, and by the mental environment. From a similar decomposition, we obtain the force on the mind as

$$f(\mathcal{B}_m, \top) = f(\mathcal{B}_m, \mathcal{B}_m) + f(\mathcal{B}_m, \mathcal{B}_p) + f(\mathcal{B}_m, \mathcal{B}_m^{me}) + f(\mathcal{B}_m, \mathcal{B}_p^{pe}). \quad (4)$$

The forces $f(\mathcal{B}_p, \top)$ and $f(\mathcal{B}_m, \top)$ constitute hybrid forces, containing components in both the physical and mental worlds. We can thus decompose them into components, for example, writing

$$f(\mathcal{B}_p, \top) = f_p(\mathcal{B}_p, \top) + f_m(\mathcal{B}_p, \top). \quad (5)$$

The deformations of the physical body and the mind depend only on these respective components in the hybrid mechanics, so motion of the physical body depends only on

$$f_p(\mathcal{B}_p, \top) = f_p(\mathcal{B}_p, \mathcal{B}_p) + f_p(\mathcal{B}_p, \mathcal{B}_m) + f_p(\mathcal{B}_p, \mathcal{B}_p^{pe}) + f_p(\mathcal{B}_p, \mathcal{B}_m^{me}). \quad (6)$$

Thus the physical motion stems from physical forces, but the origin of these forces might include mental bodies.

Many physical theories assume that self-forces vanish, that $f(\mathcal{B}, \mathcal{B}) = 0$ holds for each body \mathcal{B} . In seeking mechanical formalization of conscious agents, however, the assumption of vanishing self-forces seems undesirable. Self-conscious or self-directed action naturally suggests a mind that exerts nonzero forces on itself, though one can expect necessity of such forces in the formalization to depend on the details of the psychological organization. In particular, one might formalize a mind organized into competing sets of mental subagents (as in [33]) as involving only forces between separate mental components, so avoiding the need for nonzero self-forces.

Correlation: Few have ever thought of mind and body as operating completely independently of one another. Someone hit will hurt, and may choose to react by moving and speaking. Mental and physical events enjoy some degree of correlation. The debate in philosophy over the centuries has concerned the degree and nature of such correlation.

The constitutive assumptions that traditional mechanics uses to characterize special materials provide models for phrasing assumptions about the correlation of mental and physical events. For example, one

characterizes the rigid bodies studied in classical mechanics by a constitutive kinematic assumption, namely the assumption that the relative positions of all body parts remain constant throughout all motions. One might use similar assumptions to require that certain physical events always co-occur with certain mental events, or always co-occur in certain circumstances.

The large range of possibility between no regular co-occurrence and completely regular co-occurrence offers room for many different theories of mind. One might interpret certain identity theories of mind and some idealistic theories as positing unexceptional co-occurrence of mental and physical events. With a bit more stretching, one might interpret Penrose's [44] speculations on the role of quantum gravity in consciousness as suggesting that mental events might influence certain delicate physical events but otherwise stem directly from the major flow of physical bodily events.

When one does not assume complete correlation between mental and physical events, one expects divergence between mental and physical states to occur. We see such divergence in the experience of disoriented and insane people. The ordinary sort of disorientation represented by discrepancies between what the agent thinks about the world and the way it really is poses no philosophical difficulties. Most people lose this correlation many times a day, and in rare cases lose it seriously and permanently, but ordinarily this sort of disorientation requires only minor, even unconscious or automatic effort by the sufferer to correct the discrepancy, by changing mental or physical state to match the other. People who build robots go to some trouble to prevent such disorientation from crippling their robots by reducing thinking as much as possible to direct perception or reaction to overt physical states, sometimes employing feedback control systems to imitate mechanical models of elasticity, in which perturbation from certain states generates restorative forces to bring the configuration back to the equilibrium one. The more thinking abstracts from these lower levels to construct plans and formulate long-term desires the greater the openings provided for introducing discrepancies and disorientation.

4 Discrete mechanical systems

Broadening mechanics to cover minds and markets as well as flesh and bone requires going beyond the simple notion of hybrid mechanical systems. Making hybrid composites of ordinary mechanical systems involves no significant changes to the mathematical form of mechanics. To see mental influences as forces requires transcending mere composition to recast the mathematical axioms of mechanics in a way that separates out assumptions about smoothness of the space in which motion takes place.

Removing smoothness assumptions changes the formalism less than one might think, for in seeking to abstract away from the highly varied properties of physical materials and to separate notions in the most mathematically natural way, Noll's axioms already characterize many concepts in ways independent of the nature of the underlying space. One can see this in his development of bodies and forces. Noll's theory of bodies stays away from any assumptions about geometry, and adds geometric structure only when considering specific universes of bodies. His axioms for force systems likewise characterize a structure independent of geometry, and only at the end add assumptions about how force values relate to spatial entities. Noll [37, p. 48] himself indicates a generalization of his most abstract mechanical theory away

from familiar physical space to the level of the abstract state-spaces of general systems theory [28], but that generalization abandons all traces of mechanical concepts, and so, like hybrid systems theory, provides few insights into psychology beyond those already familiar from pure dynamical systems theory and automata theory.

4.1 Structures of discrete mechanics

We sketch how to adapt the axioms characterizing time, space, bodies, mass, and forces in individual mechanical systems so as to obtain a discrete mechanics suitable for formalizing psychology and economics. The notion of space retains its vectorial character, but with new and possibly discrete dimensions. The notion of mass changes similarly, accentuating the change in hybrid mechanics away from unidimensional masses. The modified axioms constitute proper extensions of the ordinary ones, but small changes, not radical ones. These changes leave the fundamental relationships and laws of mechanics essentially unchanged.

Time and space: Noll's axioms for space and time explicitly require that temporal instants form a continuum and that the translation space of places form a real vector space of dimension three.

We first broaden these notions by dropping Noll's axiom requiring instants to form a continuum. This allows instants to form discrete series, such as the state trajectories of automata theory, or even, should the need arise, mixtures of discrete points and isolated continua. As before, the time-lapse function induces a simultaneity relation Σ on \mathcal{W} and the corresponding partition Γ of \mathcal{W} into instants.

We retain Noll's requirement that the spatial distance functions on the instants induce a Euclidean metric making the translation space of each instant isomorphic to some vector space \mathcal{V} . We drop, however, the requirement that \mathcal{V} exhibit an isomorphism to \mathbb{R}^3 . We instead allow vector spaces over other fields of scalars, such as finite fields, and allow vector spaces of dimensions greater than three.

It may prove advantageous to weaken these requirements further. When we form the translation space of places in a hybrid mechanics, we take the product of the translation spaces of the component systems. When these spaces bear the structure of vector spaces over different fields, the resulting product has the structure of a module over a ring, not the structure of a vector space. The axioms we adopt here assume one can always decompose mechanical systems into base components in which space exhibits vectorial structure, but further experience may reveal mechanical systems in which such decompositions either do not exist or complicate the formalization excessively. In such cases, we may wish to model the translation space \mathcal{V} of some component mechanics as a module over a ring rather than a vector space over a field, perhaps providing some relation to Kalman's [28] use of modules in control theory.

Bodies: We hew to Noll's axioms in characterizing the universe of bodies, assuming that bodies form a boolean lattice of sets of body points. As in traditional mechanics, some universes will have discrete bodies, and others bodies that form continua.

While rational mechanics has discovered universes of bodies that suffice for most purposes of continuum mechanics, including the distinctive universe of Noll and Virga [41], attempts to formalize psychological structures may suggest considering still more general structures for universes. Certain models of

conscious or reflective agents seem to call for models that contain elements referring to themselves, so the acyclic subbody relation of Noll's axioms might want weakening, perhaps along the lines of Aczel's [1] non-foundational set theory. Other models may seek to draw on Scott's [48] theory of functions which makes heavy use of lattices without complements, which might motivate weakening Noll's requirement of complements to one positing only the existence of relative complements.

Motions: The notion of motion itself requires no weakening to cover discrete mechanics. The same definitions that work for traditional or for hybrid mechanics work as they stand for the discrete case.

Traditional mechanics, and hybrids over traditional mechanics, both support the standard concepts of differential calculus: limits, derivatives, integrals, and the like. These concepts either do not apply or provide less power in analyzing discrete trajectories and discrete spaces, though standard treatments of dynamics include notions of discrete derivatives and difference equations. Our development of discrete mechanics examines a variety of additional partial replacements for the concepts of analysis, including conditional vectorial analogs of differentials and tangent vectors, and "continuity" notions based on conservatism of trajectories relative to comparative similarity relations, which one can view as taking seriously Truesdell's [57] discussion of the role of smoothest-path principles in mechanics. For such reasons, we may wish to avoid assuming the existence of a distance function acting as a Euclidean metric, especially when we obtain the space of places by combining multiple spacetimes with individual metrics but no distinguished combination metric. In identifying discrete motions satisfying minimum principles, it suffices to require only that the translation space support a comparative similarity relation that tells when one translation dominates another, as in comparisons along individual dimensions, without providing an ordering of size over all possible translations.

Mass: Hybrid mechanics weakened the notion of mass to a multidimensional concept from the familiar unidimensional concept of traditional mechanics but retained the continuum character of the dimensions of mass values. For discrete mechanics we weaken this notion along with the concept of space to mass measures taking values in modules over rings rather than in $(\mathbb{R}^{0+})^n$. Indeed, mass values provide yet another motivation for regarding the translation space of places as a module, using mass values as elements of the ring of scalars to form momentum values as products of mass and spatial-translation values.

We also depart from nonrelativistic treatments of mass in permitting mass to vary with time. Traditional continuum mechanics usually avoids this need by distinguishing between bodies and the places they occupy, but textbook physics problems normally include problems in which the bodies of interest vary in mass, for example, the motion of rockets exhausting portions of their mass to obtain thrust. Time-varying mass measures play an important role in formalizing some psychological systems in mechanical terms.

Traditional mechanics formalizes the standard notion of physical mass so that finite bodies have finite mass. No general mechanical principle forces discrete mechanics to make a similar assumption. Indeed, the discrete spaces involved make formulating such an assumption somewhat problematic for masses that form infinite-dimensional modules. We nevertheless seek to hew to the physical model in this regard by supposing finiteness of values of mass in each dimension.

One might seek to extend discrete mechanics with measures of total mass that combine values in each dimension, but such measures would seem to suffer the same artificiality afflicting measures of distance that combine physical and mental positions. Measures of total mental mass might make more sense because the assumption of finite mass arises naturally in computational theories of psychology. These theories typically assume finite information content for representations and states of mind, and we can regard information content as a cross-dimensional measure of total mental mass. One might, however, find this assumption unmotivated when infinite dimensionality of mental spaces calls into doubt the supposed finiteness of the mental body.

Forces: The theory of forces in discrete mechanics retains almost all the structure given for forces in traditional and hybrid mechanics. The obvious departure comes in the character of force values. We interpret discrete force values in one of two ways. If one regards mass values as scalars in the vector space or module of translation space, then one takes force values as elements of translation space, just as in traditional mechanics. If instead mass values form a distinct type of space, one takes force values to consist of pairs of mass-translation and place-translation vectors or module elements, as this provides a structure for expressing balance of linear momentum.

The axioms of forces in discrete mechanics also depart from the continuous case in abandoning the notion of torque and rotational momentum. The standard notions of torque and rotational momentum depend in essential ways on space possessing exactly three dimensions. Discrete mechanics, not limited to three spatial dimensions, generally offers no corresponding notion, though perhaps future investigation will reveal some appropriate conceptual correspondent.

Dynamogenesis: Discrete mechanics need not take a different position on determinism, locality, and frame indifference than that taken in traditional mechanics, but the motivations for these strictures on dynamogenesis seem weaker in the discrete case.

Traditionally, many simple continuous systems appear deterministic, and recent appreciation of so-called chaotic deterministic systems has extended this expectation to seeming random behavior. At the same time, automata theory has given prominence to the notion of nondeterministic action between discrete states, and economics has emphasized the nondeterministic rationality criteria of maximizing expected utility and seeking equilibrium solutions to games. More to the point, prominent metaphysical theories have regarded people as exercising free will or the power to choose, and modern quantum mechanics regards indeterminism as a fundamental characteristic of the world. In considering such theories, axioms of determinism seem overly contentious components of a philosophically neutral mechanics, so we omit dynamogenetic determinism as a general law of discrete mechanics, and relegate it to the status of a special law characteristic of certain special systems.

Notions of locality lack clear or useful correspondents in many discrete systems, and standard quantum mechanics apparently denies it as well. Standard topologies on discrete sets provide only the discrete or indiscrete topologies, in which the notion of locality provides no useful restriction on dynamogenesis. The nontrivial discrete topologies of Scott's [48] function theory might provide more useful alternatives here.

Still other notions of locality involve non-topological notions, such as tolerance relations, and the geometric relations of cellular automata. Accordingly, we omit locality as a general restriction on dynamogenesis, as we await development of a form appropriate to covering these cases.

4.2 Illustration: reason maintenance

To illustrate discrete mechanics in a concrete setting, we summarize elements of a mechanical description of a reason maintenance system (RMS), an artificial psychological system developed in artificial intelligence [10, 18]. A RMS provides central memory, reasoning, and introspective functions to “higher-level” mental operations. The structure of a RMS reflects central organizational aspects of many artificial intelligence systems, with elements related to short-term and long-term memory, propositional deduction rules and production rules, Bayesian probabilistic networks, artificial neural networks, non-cognitive representations, and contextual linkages among representations.

4.2.1 Psychological structure and behavior

In the original conception [10], a RMS serves as an automated subsystem of a reasoning or problem-solving system external to the RMS, receiving information about inferences made and actions taken by the external system, and providing information about the current conclusions or mental state resulting from these inferences. These operations go beyond mere database entries because the RMS revises conclusions on the basis of recorded inferences, not on the basis of simple instructions to add or remove conclusions.

Representations: The RMS employs a set of structures called *nodes* to represent to itself the possible elements of computational or mental states, that is, the pieces of information of significance to the external system, such as data structure values, computational methods, or quite arbitrary aspects of mental states, such as beliefs, desires, preferences, intentions, plans, even emotions. No actual set of conclusions need contain these elements; they serve only as the reservoir of components from which the RMS constructs actual states.

Each state of the RMS represents a labeling of the set of nodes that assigns to each represented node a label of either *In* (the current state), *Out* (of the current state), or *Nyl* (not yet labeled) so that the information state of the external system consists of those state components labeled *In*. The labels *In* and *Out* do not stand for true and false, but rather indicate presence or absence in the information state presented to the external system, and beyond that, presence or absence of valid reasons supporting the presence of the information in the constructed state. We regard nodes not explicitly represented in the RMS memory as implicitly labeled *Out*. The label *Nyl* constitutes a temporary label used only during intermediate processing, and not a determiner of an external information state. Representing standard logical reasoning systems in the RMS requires representing propositions and their negations as distinct nodes, for which the separate *In/Out* labelings give rise to a four-valued logic of the sort described by Belnap [4].

Each state of the RMS also represents a secondary labeling of the set of nodes to indicate the subset of nodes for which the primary labeling represents a change from the primary labeling in the preceding state.

This secondary labeling also employs the *In* and *Out* labels to mean a node falls *In* (the set of changes) and *Out* (of the set of changes). As with the primary labeling, we regard nodes not explicitly represented in memory as implicitly labeled *Out*.

The RMS employs structures called *reasons* (sometimes also called justifications) as explicit records of inferences or changes of mental state. These reasons represent concrete inferences made by the external system. The inferences representable by reasons include nonmonotonic and unsound heuristic inferences in addition to logically valid inferences. Nonmonotonic reasons sanction conclusions on the basis of the absence as well as the presence of information, which allows one to defeat previously-stated reasons by justifying nodes formerly *Out* of the state. For example, one can use a reason of the form “Conclude Sasha can fly whenever (a) it is believed that Sasha is a bird and (b) it is not believed that Sasha cannot fly” to sanction the belief that Sasha can fly until an acceptable reason is stipulated for believing Sasha could not fly. We focus here on reason interpretations written in terms of four sets of nodes as $A \setminus B \Vdash C \setminus D$, read “*A* without *B* gives *C* without *D*”, to denote a reason bearing the interpretation that each node in the *consequence* set *C* must be *In* and each node in the *exclusion* set *D* must be *Out* whenever each node in the *antecedent* set *A* is *In* and each node in the *qualification* set *B* is *Out*. The reasons used in the original RMS used a simpler form equivalent to $A \setminus B \Vdash \{c\} \setminus \emptyset$ for some individual node *c*.

Each state of the RMS represents a set of *base* reasons. The RMS operations, described in the following, use these base reasons to construct and revise the primary node labelings.

Operations: The RMS engages in a sequence of interactions with its environment. These interactions instruct the RMS to add or remove base reasons from its store, ask after the labeling of specific nodes, and communicate changes of labeling to the environment.

When the external system instructs the RMS to add or remove base reasons, it revises the labeling of nodes to reflect the changed set of reasons. It takes note of the labeling prior to the relabeling process, and uses this to identify the nodes having changed labels at the end of the relabeling process. When labels changes, the RMS communicates this fact to its environment.

The RMS attempts to construct stable, grounded labelings from the set of recorded reasons. Stability means that each node is *In* if and only if at least one recorded reason supports it as a conclusion given the labeling of all nodes, while groundedness means, roughly speaking, that each node labeled *In* has a valid noncircular derivation from the recorded reasons. The nonmonotonicity of reasons means that adding reasons may involve relabeling some *In* nodes *Out*, and relabeling *Out* nodes *In*. This nonmonotonicity also introduces an ambiguity of interpretation that lets some sets of reasons possess more than one stable grounded labeling, some sets possess exactly one, and some sets none at all.

The RMS picks one labeling to use when at least one exists. When more than one labeling exists, the RMS attempts to pick a revision that retains as extensive a subset of the previous conclusions as possible. More generally, the RMS chooses successor revisions to minimize distance from the preceding state, with distance comparisons between transitions captured by a reflexive and transitive comparative similarity relation.

Interpretations: The organization and operation of the RMS reflects psychological and economic concerns in addition to computational considerations, in that the underlying rationale for how the RMS revises labelings involves interpreting reasons as intentions and preferences of the reasoner, and labelings as rational choices given these intentions and preferences [12, 15, 18].

Each reason expresses an intention of the reasoner about the structure or degree of coherence of its state, namely that the labeling should satisfy the condition on labelings expressed by the reason. The agent thus seeks a state satisfying all these intentions, namely a stable labeling.

Each reason in addition expresses preferences over possible labelings. For monotonic reasons (reasons with interpretations of the form $A \setminus \emptyset \Vdash C \setminus \emptyset$), these preferences rank all labelings equally. For non-monotonic reasons, the preferences rank labelings yielding the conclusions over labelings in which defeaters prevent the reason from yielding the conclusions.

Stable grounded labelings exhibit Pareto optimality with respect to these preferences. Such labelings satisfy maximal sets of nonmonotonic reasons, so that making some defeated reason undefeated requires defeating some other reason. This means one cannot increase preferability with respect to one reason without decreasing preferability with respect to some other reason, rendering the labeling an equilibrium point of optimizing processes. Indeed, this preferential import of reasons means that a version of Arrow’s theorem on the impossibility of rational social choice methods applies to RMS labeling methods as well [20].

4.2.2 Mechanical content and formalization

In the mechanical formalization of the RMS, we view the RMS and its environment as complementary bodies in the universe of bodies.

Kinematics: For uniformity, we assume nodes exist to represent each possible reason, and that we can interpret every node as a reason, even if one bearing only the trivial interpretation $\emptyset \setminus \emptyset \Vdash \emptyset \setminus \emptyset$. We write \mathcal{D} to mean the set of all nodes and reasons possible within the RMS under examination.

We take the set of possible *In/Out* labelings of nodes as the space of positions. This space forms a binary-valued vector space $\mathbb{D} = 2^{\mathcal{D}}$ over the field \mathbb{Z}_2 . Elements of \mathbb{D} thus also can represent changes to positions, with the element $x - y$ representing the change from position y to position x .

We view the set of base reasons as the mass of the RMS. Our unifying assumption means we can use \mathbb{D} as the space of possible mass values as well.

Trajectories of the RMS consist of sequences of states of node labels, label changes, and base reasons. We represent such trajectories as sequences of triples (x, m, \dot{x}) of vectors in \mathbb{D} . By construction, the label-change vector \dot{x} constitutes the velocity of the RMS, so these vector triples represent position, change, and mass vectors.

We view the velocity and mass components of such triples as constituting the instantaneous momentum $p = (m, \dot{x})$ of the RMS.

Dynamics: The most direct interpretation of RMS operations in mechanical terms considers trajectories of states more fine grained than in the ordinary external view of the RMS. A fine-grained trajectory includes

the intermediate states occupied during node relabeling as well as the equilibrium states visible to the RMS environment. Each step in this trajectory corresponds to modification of the partial labeling by application of one or more reasons. The fine-grained motion of the RMS then terminates in an equilibrium state representing a stable labeling of the nodes.

We think of external stipulations of changes to reasons as constituting forces having only mass-change components, and interpretation of reasons by relabeling operations as providing forces with only spatial (accelerative) components. In other words, we decompose the force contributions from the RMS environment and the RMS itself into an applied mass force

$$f_t^{\mathbf{a}} = (\dot{m}_t, \mathbf{0}) \quad (7)$$

produced by the environment and a spatial self force

$$f_t^{\mathbf{s}} = (\mathbf{0}, \ddot{x}_t) \quad (8)$$

produced by the RMS itself, thus yielding a total force

$$\begin{aligned} f_t &= f_t^{\mathbf{a}} + f_t^{\mathbf{s}} & (9) \\ &= (\dot{m}_t, \ddot{x}_t). & (10) \end{aligned}$$

For illustration, consider application of a single reason

$$r = A_r \parallel B_r \parallel C_r \parallel D_r \quad (11)$$

to a state σ_t , in which the reason generates a force from the state that then acts on the state to produce a new state exhibiting the appropriate velocity and position. We assume for simplicity of exposition that the reason generates only a spatial force

$$f_r(\sigma_t) = (\dot{m}_t, \ddot{x}_t) = (\mathbf{0}, \ddot{x}_t) \quad (12)$$

in which the mass flux vanishes. The RMS interprets r as a conditional prescription of node labels, in that validity of the antecedent portion ($A_r \parallel B_r$) requires application of the consequent portion ($C_r \parallel D_r$). In particular, the consequent portion represents a function from labelings to change vectors, in that if the current state contains the labeling x , the consequent $C_r \parallel D_r$ indicates a change vector given by $(C_r \setminus x) + (D_r \setminus \bar{x})$, where $\bar{x} = \mathbf{1} - x$ indicates the vector or set complementary to x . Specifically, we define the force generated by the RMS from r as providing the acceleration

$$\ddot{x}_t = \begin{cases} (C_r \setminus x_t) + (D_r \setminus \bar{x}_t) - \dot{x}_t & \text{if } A_r \subseteq x_t \subseteq \bar{B}_r \\ \dot{x}_t & \text{otherwise.} \end{cases} \quad (13)$$

This yields a velocity of

$$\dot{x}_{t+1} = \begin{cases} (C_r \setminus x_t) + (D_r \setminus \bar{x}_t) & \text{if } A_r \subseteq x_t \subseteq \bar{B}_r \\ \mathbf{0} & \text{otherwise,} \end{cases} \quad (14)$$

and so

$$x_{t+1} = \begin{cases} x_t + (C_r \setminus x_t) + (D_r \setminus \bar{x}_t) & \text{if } A_r \subseteq x_t \subseteq \bar{B}_r \\ x_t & \text{otherwise.} \end{cases} \quad (15)$$

The detailed mechanical analysis of the RMS presented in [13] also shows how one can formalize the interpretation of reasons in more interesting mathematical terms, with a set of reasons determining a potential function that determines stress tensors at given mental locations, and shows that the resulting stress satisfies Cauchy’s First Law, $\dot{p} = B + \text{div}(T)$.

Dynamogenesis: RMS behavior exhibits a nominal indeterminism that can vary with the particular implementation. Indeed, the original RMS implementation exhibited deterministic behavior by always examining reasons in the same order.

RMS behavior accords with traditional mechanics in exhibiting locality properties, albeit with a different sense of locality than the usual physical one. The notion of locality in RMS concerns the sets of antecedent and qualifier sets of nodes that form the basis of reason forces. The nodes mentioned in the reason antecedents and qualifiers form the immediate “neighborhood” of the reason. Nodes mentioned in antecedents and qualifiers of reasons having these primary antecedents and qualifiers as conclusions lie in larger neighborhoods. Indeed, some interesting RMS variations forego global groundedness constraints to operate solely within local groundedness constraints [12].

Reason forces as illustrated in the preceding exhibit a frictional character. This stems from the usual RMS behavioral requirement for motion to halt indefinitely between revisions. Other designs for reason maintenance might call for a different, non-frictional behavior.

The possibility of different implementations of reason maintenance systems reflects a diversity of psychological materials. To the extent that these implementations vary or reorder the basic operations, each embodies a different relation of dynamogenesis, and so constitutes a different type of material.

The computational organization standard in the RMS just described exhibits a degree of mechanical artificiality due to the division of motion into portions in which the environment acts on the RMS and the portions in which the RMS adjusts its internal state in response to those external impulses. Although church bells and other physical systems possess such an organization, nothing in ordinary mechanics enforces such a partition of motion. It might prove fruitful to explore alternative RMS organizations in which external impulses can interrupt or intermingle with internal state adjustments. Such organizations might prove more appropriate for distributed systems. The rational distributed reason maintenance system described in [19] exhibits something of this character.

5 Mechanical rationality

The mathematical theory of economics and the theory of computation today represent the major successes of mathematical analysis outside of the physical sciences, each bringing a comprehensive array of mathematical concepts to bear in constructing theories of great depth. Economics conceives of rationality as the property

of acting optimally with coherent beliefs and preferences or utility judgments. Informatics conceives of machine computation as finite compositions of finite operations on finite representations. The modern theory of artificial intelligence seeks in part to capitalize on the successes of these fields by combining their concepts to construct rational machines.

We now turn to examine some implications of mechanical formalizations of mind for understanding the economic notion of rationality, and in particular, notions of rationality reflecting limits on reasoning and judgmental powers observed in human behavior.

5.1 Structure of economic rationality

Economics develops the theory of rational action in stages. The first stage characterizes the pure notion of rational action in terms of choosing maximally preferred actions from the actions available at some time [62]. The second stage extends the first by characterizing preferences among actions in terms of preference among the states resulting from actions and the probabilities of these results obtaining [47]. A third stage extends the second by characterizing result preferences and probabilities in terms of preference and probability information concerning properties or characteristics of result states.

We summarize here the concepts involved in this development, employing versions of the concepts chosen for the common usage they enjoy without intending to imply they enjoy greater correctness than the many alternative conceptions proposed by economists and philosophers. The discussion of limitations on rationality that follows examines some important dimensions of variation for the standard conceptions.

Rational choice: The axioms for rational economic agents characterize a notion of preference-based choice among alternative actions. The theory starts with a set $\mathcal{A} = \{A_1, A_2, \dots\}$ of alternatives and a binary choice relation $\succsim_{\mathcal{A}} = \succsim$ of *weak preference* over \mathcal{A} , where $A \succsim B$ means that the agent weakly prefers B to A in that it finds B at least as desirable as A . The theory defines two additional relations in terms of weak preference. The theory defines *indifference* among alternatives, written $A \sim B$, so that $A \succsim B$ and $B \succsim A$, meaning that the agent finds whatever differences exist between the alternatives leave them equally desirable. The theory defines *strict preference*, written $A \prec B$, so that $A \succsim B$ but $A \not\prec B$, meaning the agent finds B more desirable than A .

Three axioms characterize the notion of ideal preference by requiring that weak preference constitute a complete preorder, that is, a complete reflexive and transitive relation. These requirements entail that strict preference is transitive and asymmetric, and that indifference is an equivalence relation. A fourth axiom then requires that the agent choose alternatives maximal with respect to the ordering of weak preference. The axioms on preferences ensure that every finite set of alternatives offers at least one rational choice. Infinite sets of alternatives can lack any rational choice without constraints on the nature of the preferences and the structure of alternatives.

Expected utility: The *subjective Bayesian* theory of rational decision adds the notion of a set \mathcal{S} of states that can result from actions and in which the agent takes action. It posits another preference order $\succsim_{\mathcal{S}}$ over

these states that compares desirability of states rather than actions. It posits a belief function $p_A : 2^{\mathcal{S}} \rightarrow \mathbb{R}$ for each action $A \in \mathcal{A}$ that assigns to sets of states the agent’s degree of belief that states in this set result from A .

Three axioms require $\succsim_{\mathcal{S}}$ to constitute a complete preorder, just as with $\succsim_{\mathcal{A}}$. These axioms imply the existence of numerical *utility* functions $u : \mathcal{S} \rightarrow \mathbb{R}$ that represent $\succsim_{\mathcal{S}}$ in the sense that $A \succsim_{\mathcal{S}} B$ iff $u(A) \leq u(B)$. Another axiom requires that the belief functions representing beliefs about the consequences of choices constitute probability measures over result states. One can develop this probabilistic axiom in turn from further axioms about belief comparisons, but we do not summarize those here (see [47]).

The theory combines the numerical utility and probability functions into an *expected utility* function $\hat{u} : \mathcal{S} \rightarrow \mathbb{R}$ such that $\hat{u}(A) = \int_{\mathcal{S}} p_A(S)u(S)$. A fourth axiom then makes preferences among actions correspond to comparisons of expected utility by requiring that $A \succsim_{\mathcal{A}} B$ iff $\hat{u}(A) \leq \hat{u}(B)$.

Multiattribute decision making: The theory of multiattribute decision making extends the concepts of subjective Bayesian decision theory by decomposing utility functions over states into functions over properties or characteristics of those states. The set of properties of interest induces a multidimensional representation of the set of states. The theory then studies utility functions over the multidimensional representation that can take the form of functional compositions of “subutility” functions over lesser-dimensional subspaces of the multiattribute representation.

Recent approaches to machine computation of rational decisions augment multidimensional decompositions of preference and utility information with multidimensional decompositions of probability measures, for example, using Bayesian networks to express probability distributions relatively succinctly in terms of causal relations among propositions holding true in different states [43, 6].

Group decisions: The theory of rational group decisions extends the theory of rational individual decisions just summarized by setting requirements on choices made by a group of rational individuals [3]. The theory of rational group decisions consists of three axioms characterizing the group preference order $\succsim_{\mathcal{G}}$ in terms of member preference orders $\{\succsim_i \mid i \in \mathcal{G}\}$. The first axiom of collective rationality requires that the group preferences derive from a function over all possible rational preference relations of the individuals. The second axiom, of Pareto optimality or unanimity, insists that the group preferences agree with the individual preferences for alternatives on which the individuals agree. The third axiom, of “independence of irrelevant alternatives”, requires that the ranking of two alternatives in the global order depends only on how the individual orders rank those two alternatives, independent of how they rank other alternatives relative to the given two.

Social fairness concerns in human decision making suggest adding a fourth axiom of non-dictatorship that ensures none of the group members acts as a “dictator” whose preferences automatically determine the group’s, independent of the other individual orderings. Arrow [3] proved this requirement conflicts with the first three. Some theories of group rationality abandon or modify some of the first three in order to adopt this fourth restriction, but the first three suffice for our purposes here as they apply to cases of nonhuman decision making in which fairness concerns need not prove essential.

Markets represent a familiar context for group decisions. The competing individual preferences determine demand functions. Markets provide auctions in which these group demand functions combine with supply levels to produce equilibrium prices or exchange ratios for sets of goods.

5.2 Limitations on rationality

Economic rationality represents an ideal, not a practical observation. While no one expects reality to fit the ideal perfectly, observed individual behavior approximates the ideal conception of rationality only poorly. Psychology finds the economic conception of rationality suspect for several reasons, including the amount and precision of information it attributes to the agent, the rapidity it presumes in updates to this information, and the way the economic conception dismisses or ignores intentions, habits, and other shapers of human choices. Some consider these faults severe enough and common enough to call into question the suitability of the ideal as the standard by which to judge rationality.

Simon [52, 53] presented one of the first critiques of the ideal conception of economic rationality in his theory of bounded rationality, in which he sought to take seriously the limits on ratiocination, information, and consistency common to people. Simon pointed out that although people enjoy seemingly endless possibilities for action at every instant, they can comprehend only a few samples or dimensions of variation in these possibilities. Embodying the ideal theory in a world as complex as ours apparently requires fairly complete probabilistic and preferential information about truly vast or even infinite numbers of possible circumstances and events. Economic rationality encompasses large action spaces without difficulty, but one simply does not see this sort of synoptic awareness of the possibilities in people, who routinely find themselves or others ignorant of the consequences of their actions, and find making decisions overwhelming when faced with the myriad possibilities and ramifications of even simple decisions. Similarly, an agent unable to comprehend sufficient quantities of information will experience difficulty observing the inconsistency of separate pieces of information. Decision-making based on inconsistent preference information can result in failure to find any rational choice when inconsistent preferences rule each alternative strictly worse than others.

Economics shows how with special effort one can hide some of the complexity of real circumstances of choice by circumscribing the scope of each decision to a carefully constructed “small world”. In spite of such practices, the issue of the quantity of information required by rational agents becomes inescapable in subjective Bayesian decision theory, which requires agents to employ Bayesian conditionalization to update beliefs to reflect new information. Conditionalization implies that the beliefs of the agent at any instant suffice to determine the appropriate updates for every possible sequence of future information updates, leading quickly to belief states representing infinite quantities of information.

In the following, we interpret aspects of these psychological limitations in terms of mechanical characteristics of mass, force, and effort. Some psychological limitations involving update reflect the role of mental mass; some limitations involving habit and character reflect constitutional properties of the psychological material in question; some limitations of attention reflect the additivity of forces and energetic properties; and some limitations of consistency reflect multiple bodies as well as the additivity of forces. Roughly speaking, one can see masses in memory, intention, and some habits; forces in motivations, desires, sensa-

tion and attention; energy in effort; and constitution in character.

The discussion here represents less formal and more speculative ideas than those treated in the preceding.

5.2.1 Kinematic limitations

Economics characterizes the states of a rational agent in terms of rational preference and belief function over alternatives and outcomes. Psychologies for agents of limited rationality and more general psychologies as well expand the content of mental states and alter the constitutive assumptions of completeness, consistency, and competence provided in the axioms of economic rationality.

Space: One constructs the space of positions in the theory of ideal rationality from the sets of all preference orders over alternatives and all belief functions (probability distributions) over alternatives and outcomes. We might formalize the space of mental positions of an ideal rational agent in terms of an inhomogeneous module consisting of a product $U \times P$ of the vector space U of all utility functions over outcomes, using \mathbb{R} as the field of scalars and pointwise addition as the vector sum, with the vector space P of all probability distributions over actions and outcomes, using \mathbb{Z}_2 as the field of scalars and the normalized sum of distributions as the vector sum.

Many economic analyses use an even simpler conception of position, that of just the belief functions, and treat the preferences of the agent as fixed in its constitution, with all variation of decision obtained through variation in the agent's beliefs. While one can embed many desired ranges of agent behaviors into such belief-centered representations, in modeling ordinary human cognition one often finds it more natural to think of the agent changing preferences as well as beliefs. Economists tend to avoid models involving change of preference mainly because they possess no empirical or prescriptive theory of how preferences change. One can formulate such changes in the standard framework of rational action by taking the self-management perspective exemplified by Pascal's wager and James' will to believe [42, 27, 11, 16], in which the agent makes choices about adopting or abandoning discrete attitudes or sets of attitudes. The psychological literature provides other patterns for preference revision as well [26, 50, 45, 63, 17].

The self-management perspective of ordinary psychology implies more structure to states of belief and preference than one can recover from utility functions and simple probability distributions. In consequence, the two-component space of positions seems unsuitable for a theory of bounded rationality encompassing discrete actions and bearing close similarity to psychological theories based on discrete or semi-discrete conceptions of beliefs, desires, intentions, and the like, or even the structured representations used in computational decision theory.

At a more fundamental level, when one characterizes mental positions directly in terms of utility functions and probability distributions, one builds the axioms of economic rationality into the very fabric of mental space. This makes weakening the notion of rationality fairly difficult, as one must work against the ideal nature of the points of space themselves.

To capture discrete psychological theories more directly, and to make weakening the rationality assumptions easier, we decompose mental positions into discrete elements of preference and belief information. Specifically, we replace the inhomogeneous two-component space of ideal rationality with homogeneous

multi-dimensional binary vector spaces, as exemplified by the mechanical formalization of the RMS. To capture wider ranges of psychologies, we may go further and employ inhomogeneous modules that augment the basic attitudes and representations of rational economic agents with other attitudes, feelings, emotions, sentiments, perceptions, and sensations.

Configuration: In decomposing the space of positions into discrete attitudes, we move the consistency and completeness axioms of rationality out of the characterization of the structure of space, where they stand unmodifiable in particular mechanical systems, and into constitutive assumptions about the structure of configurations and dynamogenesis, where we may modify them freely to characterize agents of different degrees of rationality. Constitutive assumptions of consistency and completeness in restricted cases exhibit a purely kinematic character akin to more familiar kinematic constitutive assumptions for rigid bodies and crystalline materials.

Weakening upper bounds: To acknowledge Simon’s critique of the informational demands of ideal rationality, we drop the completeness assumptions underlying the ideal theory of rational choice and consider agents with partial, not total, weak preference orders. We similarly give up consistency of the preferential order and allow the agent to possess conflicting pieces of preference information.

With inconsistent information embodied in separate representations, it becomes sensible to treat the construction of the mental state of the agent as a group decision problem, subject to an appropriate adaptation of the axioms of group rationality. We noted this approach earlier in commenting on the preferential interpretation of RMS reasons and the Pareto optimality of RMS conclusion states.

In light of the weakening of the conditions of individual rationality, we weaken the theory of group decisions to require only partial preferences of the group members and of the group. These weakenings of assumptions about preferential completeness motivate adoption of an additional axiom of group decision making, a rule of “conflict resolution” that requires that the group preferences relate two alternatives whenever any group members do, as in [20].

Imposing lower bounds: Though we allow beliefs and preferences to exhibit incompleteness and inconsistency, abandoning the ideal rationality requirements for these attitudes leaves open a broad range of specialized forms of completeness and consistency. Indeed, one can capture restricted forms of completeness and consistency using the formalization of *information systems* developed by Scott [49] to understand computation over partial representations in terms of limited consistency and completeness conditions. To express constitutive assumptions in this way, we formulate the kinematic states of the agent as the sets of elements of the information system closed with respect to the limited completeness relations and consistent with respect to the limited consistency relation. We use these to characterize the native intelligence of an agent in terms of lower bounds on coherence, consistency, rationality, and knowledge [12, 15, 18]. One can expect stipulations of constitutive consistency and completeness for restricted classes of other mental attitudes and operations to form key elements of specifications for artificial psychologies and agents.

People have recognized constitutional characterizations of degrees of rationality longer than they have recognized a formal or informal subject called psychology. Geniuses and dullards alike sometimes err in

allocating effort and take stupid actions, but the less able suffer from abnormally low lower bounds on their constitutional abilities. They can perform the same reasoning in principle, but must perform much more of their reasoning consciously, at enormous cost in attention and resources. Because they must attend to the bookkeeping needed to keep attention focused as well as to the external focus of reasoning, the half-witted find things not uniformly twice as hard, but exponentially (in the complexity of the problem) harder than their full-witted cousins. Even extreme diligence finds this handicap hard to conquer. The difference between novices and experts constitutes an important special case of this difference personally familiar to most people. Novices, even when possessed of adequate instructions, must perform every step consciously, and expend much effort in keeping track of their progress and of their understanding of the instructions. The expert performs almost all of this reasoning automatically and seemingly effortlessly. Normal novices have adequate automatic reasoning powers, but have not yet committed their instructions to these powers. Their intelligence in other arenas helps but little in their new subject.

5.2.2 Dynamic limitations

Economics characterizes the behavior of rational economic agents in terms of acting to maximize preferability or expected utility and in terms of using Bayesian conditionalization to update beliefs on taking action or receiving new information. Psychologies for agents of limited rationality modify these characterizations of decision making and belief change to acknowledge impediments to change, bounds on forces, the complexity of deliberation, and the operation of habitual behaviors.

Ordinary mechanics limits motion by requiring that bodily motions exhibit inertia linearly related in magnitude and direction to applied forces. Ideal rationality lacks such bounds on motion because arbitrarily small changes can produce arbitrarily large effects arbitrarily rapidly. For example, Bayesian updates can wreak large changes in the belief state of an ideal reasoner, as when someone reliable tells the reasoner that everything the reasoner knows is wrong. The ideal theory presumes that even such large belief updates take place prior to decisions about the next action, so as a practical matter one views the ideal theory as assuming arbitrarily large updates take place arbitrarily quickly. Non-ideal humans, in contrast, change their mental state relatively slowly and with difficulty. We can interpret some of these limitations on rationality in terms of limitations on forces and the existence of mental inertia.

Forces: The forces that move bodies have both magnitude and direction. We can interpret limitations on the magnitude of forces as limiting the rate at which the environment stimulates change in the agent's mental state, and limitations on the direction of forces as influencing the coherence and focus of the agent's actions. Combining forces with the changes they produce yields notions of work related to familiar economic measures.

Magnitude: Traditional mechanics divides forces into contact and body forces, with contact forces acting across a shared boundary and body forces acting directly on a body without need for contact. In psychology, one naturally views communications between agent and environment, and between agent interior and periphery in terms of contact forces. Indeed, the role of the shared boundaries in agent communication with

senses, effectors, and its environment provides a natural connection between mechanical and psychological concepts. In addition, although one might view desires of the agent as exerting body forces on the agent, one might instead regard these self-forces as contact forces operating on the agent as a whole (its boundary with itself).

As summarized earlier, traditional mechanics assumes that body and contact forces vary smoothly as one varies the size of the contact area and body mass and volume under examination. More precisely, it assumes that contact forces vary continuously with the area of the contact boundary and that body forces vary continuously with both the surface area and mass of the body. One formalizes these smoothness conditions in terms of inequalities that bound contact forces by some multiple of the area of the contact surface and that bound body forces by multiples of the surface area and mass of the body, with the body mass bounded in turn by a multiple of the body volume.

Such area- and mass-proportionate bounds on the magnitudes of forces, so familiar in physical materials, also appear in the mental realm. The first fundamental theorem of Shannon's [51] theory of communication provides limits on the number of bits transmissible across a noisy channel. In the psychological setting, we may view the communication boundary bodies as finite-area contact regions, and view the dimensionality of the state spaces of the communications boundaries as imposing such limits on contact forces, with the dimensions (orthogonal components) of the state space proportional to the "dimensions" (size) of the boundary body.

While bounds on contact forces and channel capacities seem closely related, one finds in psychology no immediately available correspondent for bounds on body forces, although some theorems of distributed computation might bear connections to these limits. One person can tell another more than he knows because the other person may be able to infer things from the statement that the first one does not know. Indeed, exchange of information may lead to the people both inferring things known by neither of them at the start. If mass measures relate to measures of knowledge, the magnitudes of distributed knowledge might thus suffer bounds expressed in terms of the masses (knowledge) of the interacting agents. One can of course apply the same reasoning as for contact forces and see that the dimensionality of the mental momentum space limits the magnitude of body forces on the mind. In this case, we interpret the dimensionality of the momentum space as mass and volume rather than as an area.

Direction: Attention, as such, does not exist in the ideal theory of rational action. The fickle agent can jump from one activity to another to track every minute shift in the levels of expected utility of alternative activities. Although such behavior might serve the agent's purposes well, shifting activities can consume significant cognitive and external resources. Such task-shifting costs normally slow progress on intended activities and thus represent costs worth avoiding. Realistic rationality requires that the agent account for these costs in assessing expected utility. Attention thus appears only implicitly, as an epiphenomenon of the structure of expected utility assessments.

We can use mechanics to understand issues of attention more directly in terms of resultant forces arising from a superposition of competing forces. In informal terms, attending to something means focusing one's perception and action onto the object of one's attention. When forces serve to embody the motives of action, attention means shaping these forces in a specific direction. Failing to shape forces in that direction means

that forces in other directions can provide the motives for action. These other forces represent distractions from the intended focus of attention. Agents that seek to maintain a focus of attention thus either must move to avoid circumstances that would produce distracting forces, or must generate restorative forces to counteract or defeat the forces of distraction. This latter approach reflects the character of elastic materials, discussed further in the following.

Mechanical forces can exist even without motion when the balance of forces on a body places it in an equilibrium configuration. One can think of static equilibrium states, as seen earlier in the formalization of the RMS, as constituting a special case of the directionality of attention, that of a directionless zero resultant force rather than a nonzero resultant force indicating a unique direction.

Work: Simon pointed out long ago that ideal rationality lacks measures of the effort required to make decisions. The modern theory of computational complexity studies a number of specific and abstract measures of procedural difficulty, some of them inspired by the concepts of physical science [23]. Recent theories of rational allocation of inferential resources have made progress addressing questions about mental effort. The mechanical perspective brings yet another way of thinking about these issues by interpreting effort not just as assignments of resources to tasks but as mechanical work, the exercise of force across distance. Resources devoted to (or consumed by) tasks amount to changes of resources (mass) accompanied by changes of position (task achievement), which forms a quantity of work. One can interpret some economic measures of costs and benefits as work in a similar way.

Inertia: We formalize the notion of mental inertia for rational agents as in the preceding discrete mechanical formalization of the RMS. We identify one portion of the mental state as the mental mass, and another portion as the mental position. Mental forces then act to change masses and velocities.

The difficulty and slowness with which real agents change their mental state constitutes the most evident limitations on rationality after the limitations on completeness and consistency noted earlier. Common observations of these limitations often use mechanical terms like inertia and force. Quite apart from the common usage of these terms in both informal and formal language, one can see reflections of the mechanical connection between momentum and force in truisms of popular psychology. Consider first “the more you need to change, the more you have to force yourself.” We can read this truism as stating a monotonicity relation between the size of changes and the size of the required forces and work done. Of course, the truism does not imply the mechanical notions, but it stands consistent with them. Consider second “the more you know, the harder it is to change your mind.” We can read this truism as stating a monotonicity relation between the mass and the force required for given changes.

As these examples suggest, an informal mechanical interpretation of thinking naturally relates slowness of change to inertia. From the same perspective, the unreality of ideal rationality appears in the large changes from small impulses characteristic of actions determined by maxima of an expected utility field generated by instantaneous beliefs and desires. This lack of proportionality between the new information causing changes and the resulting changes suggests we view these systems as operating without inertia.

In contrast to ideal rationality, realistic psychologies involve many elements that exhibit degrees of persistence suggestive of the persistence of motion due to inertia. First, humans and artificial agents often

use slower-changing intentions and plans to mediate action, with beliefs and desires guiding decisions to adopt or abandon intentions that persist until the agent acts on them or makes further decisions that change them.

Second, almost every psychological theory involves some notion of memory. The preceding mechanical interpretation of the RMS provides an example, identifying RMS memory as the set of base reasons. Similarly, cognitive architectures based on production systems employ persistent rule memories intended to model long-term memory. Production systems also involve a notion of short-term memory for small sets of conscious elements, but we might better view these as positional elements, much as we viewed the constructions from base reasons in the mechanical interpretation of the RMS.

Third, decision-theoretic beliefs and desires share a quality of entrenchment with other aspects of mental states. Several theories of belief revision employ a notion of “epistemic entrenchment” [22] that orders beliefs by degree of entrenchment, such that when the agent removes beliefs to retain consistency, it removes the least entrenched beliefs first. Similar notions apply to changes to preferences and other mental qualities as well [17].

Fourth, habit shapes much of human action, and corresponding persistent tendencies or patterns of behavior also appear in artificial agents. One can interpret the rules of production rules, the behaviors of subsumption architectures, and the reasons of the RMS in such terms. Indeed, our psychological interpretation of RMS behavior involves both intention and habit by interpreting reasons in the base reason memory as representing intentions that express certain habits of reasoning. Mechanically, we can view the store of habits as part of the mass of the agent alongside purely passive memories. The action of habits forms part of dynamogenesis in the agent, as discussed in the following.

Dynamogenesis: Although traditional mechanics provides numerous theories of specific types of materials that give rise to different patterns of forces, the typical uniformity of physical materials that ensures that models of one sample of rubber or iron apply as well to the next batch does not hold in psychology at familiar levels of abstraction. People come in batches of one, with each individual batch extremely different in details, if not in overall form, from the next. In spite of this, everyone recognizes certain common types of psychological character that give insight and enable reasonable, though not perfect, predictions of individual behavior. We thus reconceive theories of these character types as theories of special psychological materials, with each special theory characterized by special constitutive laws of the material.

The characteristics of materials related to dynamogenesis correspond to the notion of personal character in psychology and legal, organizational, or economic character in economics. In psychology, personal character, whether the character of a specific individual or the character of a personality type, provides the structure of underlying motivations and behavior, as distinguished from the instantaneous impulses and thoughts of the person. In economics, attention lies more on multiagent systems, which we may view as extended bodies, than on individual agents. Indeed, mathematical economists have developed theories of special economic types including n -person games, deterministic games, and markets with a continuum of agents. More generally, economic character reflects organizational goals, such as missions or profit maximization, and behavioral rules, such as negotiation and communication protocols, social norms, or

legal constitutions.

As noted earlier, mechanics distinguishes laws relating forces existing at an instant to motion at that instant from dynamogenetic characteristics of materials that give rise to the instantaneous force values. Psychology does less well at making this distinction clear in language. For example, terms like desire and motivation can refer to specific temporal attitudes and sentiments that act at some instant as well as to underlying desires and motivations that persist and on occasion give rise to specific temporal manifestations. In the mechanical view, we interpret desires both as forces and as generators of forces. We interpret individual desires at particular instants as particular forces acting on or within the agent, and desires more generally as mental elements participating in the genesis of particular forces. The more general type of desire or motivation represents either part of the mass of the agent or an aspect of the agent's character or constitution.

The generation of instantaneous forces in human psychology divides roughly into deliberate and habitual dynamogenesis. Deliberate dynamogenesis yields forces as the result of deliberation, generally a conscious process guided by many subsidiary choices and forces involved in coming to decision. Deliberate dynamogenesis involves intentional forces that seek to maintain attention or focus in the agent's actions and state. Habitual dynamogenesis yields forces through the action of unintentional habitual behaviors, which generally operate unconsciously even if recognized consciously by the agent. Habitual behavior, operating "unthinkingly" and independently of current beliefs and desires, can result in the agent taking certain actions he judges irrational immediately afterward.

We briefly sketch some patterns of dynamogenesis that reflect abstract forms familiar in traditional mechanics: fields of attraction and repulsion, and materials exhibiting elasticity and plasticity.

Attraction and repulsion: Familiar physics includes a variety of laws that derive attractive and repulsive forces from different types of charges, such as electromagnetism and gravitation. A variety of psychological and economic theories also involve sources for attractive and repulsive forces.

The economic theory of preference and utility provides perhaps the simplest example of such forces, with the utility surface over alternatives resembling a potential function for which gradients indicate directions of attractions to local maxima and repulsions from local minima. The axiomatic simplicity of economic theory results from reducing all psychological motivations to the preferences they entail concerning possible choices made by the agent.

Most psychological theories, in contrast, discern a variety of underlying motivations, emotions, and sentiments that we can interpret as generating attractive and repulsive forces. Unlike the typical charge-generated forces of physics, the forces arising from psychological elements can exhibit conditionality, generating zero force in one set of circumstances while generating nonzero force in another, as was exemplified in the preceding by the interpretation of RMS reasons. This conditionality of some types of forces widens the scope of dynamogenesis from simple superposition to competition and defeat of one force-type by another.

Goals chosen deliberately constitute the most familiar class of sources of attractive and repulsive forces in deliberate dynamogenesis. Underlying or constitutional motivations, desires, dislikes, fears, hates, and many other psychological elements generate attractive and repulsive forces in habitual dynamogenesis. These elements provide the constant influences acting on the agent in every deliberation, influences the

agent can exploit when they align with its intentions, and influences the agent must fight when they diverge from its intended path. Shand [50] characterizes many types of emotions and sentiments in terms of particular patterns of forces acting to increase certain types of order in mental states. From his perspective, one might think of each distinct type of emotion or sentiment as representing a different type of charge, each generating a different characteristic field of position-dependent forces that superpose to yield the typically confused forces felt by humans.

Elasticity: In an elastic material, deformation from a relaxed configuration generates a restorative force. Hooke's law for small deformations of springs provides the most familiar instance, with the restorative force proportional to the magnitude of the deformation. We may interpret some forms of both intentional and unintentional psychological behavior as exhibiting an elastic character.

The deliberate form of elastic forces arise when the agent generates and applies forces to maintain focus by counterbalancing the forces of distraction, leaving only forces in the intended direction. In such cases, the intended direction of motion constitutes the zero-point of the elastic response. Deliberative agents expend significant effort in maintaining intentions and plans. When new information undermines the motivations for some intentions or portions of plans, the agent might choose to reconsider these. When new stimuli motivate new plans that threaten to draw attention from primary concerns, the agent might choose to temporarily defeat or postpone the new plans to keep action focussed on the primary concerns. More generally, one may interpret many forms of negative feedback in terms of elastic response. Indeed, many simple control systems employ negative linear feedback modeled directly on Hookean elasticity.

Habitual forms of elastic forces can arise through the action of habits that act to maintain habitual positions. Such habits can either further or impede the progress of intentional actions. As everyone knows who has fought bad habits, these habits tend to generate fixed desires seemingly independent of intentional consideration, desires that can undermine resultant forces in a way that creates weakness of will or backsliding. Indeed, applying course-correcting forces to counteract these habitual forces does work, accounting for the effort so familiar in maintaining attention. Because habits tend to restore the distractions over and over, seeking to maintain a course of action in spite of recurring habitual distractions involves a great deal of work. When strong desires or fears threaten to undermine the ability to act on intentions, the agent might choose to set up additional intentions as backups or to take preliminary actions to shape the circumstances in ways that prevent backsliding or weakness of will.

Plasticity: In perfectly elastic materials, removing a deforming force permits the elastic force to restore the material to the undeformed configuration. In plastic materials, deformation may produce a restorative force, but this restorative force need not suffice to return the body to its original configuration on removal of the deforming force. We may also interpret both deliberate and habitual forms of psychological behavior as exhibiting a plastic character.

The most common manifestation plasticity appears in the typical and habitual conservatism of human thought when the agent adopts or abandons elements of mental states. Intended changes in mental states produce some of the differences between the old state and the new one, but other changes may accompany the intended change as side effects incurred in moving to a new relaxed or equilibrium state of the mental plastic material.

One sees this in theories of belief revision, and indeed, in the action of the RMS, in which abandoning an assumption yields a new state satisfying constitutional consistency and completeness requirements but standing as close as possible to the previous state. One characterizes the notion of conservative revision in terms of a distance notion for mental states, or more generally a comparative similarity relation that embodies a preorder on translation vectors. Even such a weak comparison suffices to identify the minimal changes that include the intended changes while satisfying constitutional requirements.

Habitual plasticity forms an important concept in discrete mechanics. It engenders a conception of smoothest paths that helps replace in part the notion of continuous trajectory in traditional mechanics. Here motion occurs through small changes, as in Galileo's mechanics. It operates by replacements of discrete chunks that obey minimality principles reminiscent of the least-change principles of variational mechanics.

One finds a common manifestation of deliberate plasticity in familiar learning and exercise strategies, in which the agent pushes himself beyond current abilities in order to expand those abilities and come to rest at some new equilibrium point beyond the current frontier. Some negotiation protocols exploit a similar form of plasticity, making knowingly-unacceptable demands or bids intended to reset the expectations and appetites of other participants.

6 Discussion

The preceding sketches the beginnings of a path toward understanding psychology and economics in mechanical terms, but hardly answers all the questions that arise. The mechanical perspective differs in significant ways from the usual one brought to psychology and economics, ways that complicate appreciation of this new point of view. This section addresses several topics we hope aid in understanding issues of mechanizing psychology and economics.

6.1 Defining mechanics

Standing on the post-Newtonian discoveries of non-Euclidean geometries, Hilbert emphasized that mathematical notions have no definite meaning apart from the axioms that characterize their relations to one another, and called for the axiomatization of physical science in the sixth problem he gave in his famous address of 1900 [25]. We recognize that we get axioms for a subject by abstracting from specific examples of known interest, and by refining the axioms after looking to see what satisfies them that should not. Once formalization of disparate examples forces the axioms to exhibit some level of abstraction or generality, we generally should regard additional satisfiers as new examples of that concept, not as uninvited guests.

The mechanical axioms devised by Noll to solve Hilbert's sixth problem really do seem to capture the general intuitions underlying the subject of mechanics. They identify the basic structure and interrelations of the concepts of bodies, masses, motions, and forces. The level of generality they exhibit stems directly from the variety of acknowledged mechanical systems Noll intends them to characterize. We can rejoice in discovering that these axioms also characterize important psychological and economic systems with only minor and inessential changes, changes mainly aimed at removing assumptions of continuity rather than at

changing any of the directly mechanical characteristics.

Mechanics in its essence involves no notion of continuity. People understood motion informally long before mathematicians developed any formal notion of continuity. To see this, recall the difficulties Zeno revealed in trying to understand continuous motion before adequate intellectual tools had been invented, and that Galileo explicitly treated mechanical motion as concatenation of many small discrete changes. Indeed, much of what we call mechanical today involves discrete motions, gears and clockworks, kinematic operations intelligible quite apart from the continuous mathematics needed to understand smooth motions. Extending mechanics to discrete systems thus seems a natural enterprise, not like abandonment of some tenet essential to the identity of the subject. Force and motion constitute the central notions of mechanics. Neither of these notions necessarily involves continuity.

6.2 Prediction and design

Traditional applications of mechanics use mechanical laws to make predictions by two means: by applying theorems that yield conclusions about future behavior from limited facts about past and present behavior, and by detailed numerical calculation or simulation of behavior. In practice, scientists and engineers address most problems of prediction with numerical calculation, sometimes due to ignorance of or impatience with theoretical conclusions, but more commonly because theoretical methods do not seem to provide a solution. In fact, people often turn to simulation because not all predictions prove amenable to analytic solution, even in traditional mechanical applications. If one cannot predict the behavior or underlying structure of a system analytically from its axioms, at least one can try to simulate the temporal evolution and look to see what happens. In most of science and engineering this means solving differential equations numerically.

The discrete character of much of the structure of psychological and economic behavior removes many prediction problems from the realm of applicability of traditional differential equations and numerical calculation. This does not make the notion of simulation less important. In the discrete realm of psychology, simulation amounts to application of discrete rules or transition systems, such as application of argument steps to obtain conclusions or reasoned changes of mental state. This sort of simulation is exactly that practiced in artificial intelligence and cognitive simulation, where one writes a program to describe (rather than compute) the desired behavior and then runs the program to observe the resulting behavior [14]. Though different in the operations performed, symbolic or reasoned calculation corresponds directly to ordinary numerical integration, which one may view as a method for discrete simulation of continuous flows.

The mechanical perspective on psychology presented here may lead eventually to a practice of artificial intelligence less reliant on simulation, but the primitive state of mathematics appropriate to psychology and economics promises to impede rapid change in the near future, because the ability to make powerful theoretical predictions that short-circuit or simplify simulations depends critically on the power of the available mathematics to formalize and analyze the central structures under study.

Even if simulation remains the rule, one can expect mental mechanics to provide an increase in the degree of clarity and intelligibility of the designs proposed and examined by artificial intelligence. Natural sciences seeking to understand some phenomenon must work with the situation they find, but engineers of artificial systems can choose designs to facilitate prediction and analysis as well as to minimize cost or to

facilitate manufacture, and so render the limitations of theoretical prediction less onerous in engineering than in science.

In addition, the mechanical concepts of mass, force, and the like provide a new theoretical language with which to specify desired behaviors, initial conditions, and the material laws embodied in designs. One might expect conformance of system behavior with human predictions based on informal mechanical truisms to represent one of the more important uses of these concepts. Such uses aim not to facilitate predictions by the design engineer, but to improve the accuracy of predictions made by the users of the designed system.

6.3 Open problems

Discrete and hybrid mechanics offer many open problems for investigation: some mathematical, some theoretical, some experimental.

Mathematical: As mentioned in the discussion of prediction, current mathematics lacks many of the tools one might want in seeking to understand the novel structures formalizable in discrete and hybrid mechanical systems. Although exceedingly rich and deep, the mathematics of today provides the most leverage on the oldest problems of continuous geometry and motion. Mathematicians have made some inroads on providing tools for analyzing discrete systems. Scott's [48] theory of computation in terms of continuous functions on lattices, various theories of integration over lattices, studies of symbolic dynamics, various new logics; all these offer hopes and leads, but one can expect the needed concepts to take some time to develop. We hope that having mechanical formulations of new psychological and economic materials will spur mathematical investigation today as much as mechanical ideas motivated the development of classical mathematics.

Theoretical: Although we indicated how one might identify certain concepts of certain psychological theories as instances of mechanical concepts, much theoretical work remains to understand which psychological concepts conform to mechanical requirements. In addition, discrete mechanics lacks at least one concept playing a role in continuum mechanics, namely the concept of torque. Mathematical progress may develop concepts corresponding to cross products to provide notions of rotational momentum and torque in discrete spaces, but we need not wait to see the outcome of such investigations to look for additional concepts that provide analytical and predictive power in discrete systems.

Other directions for investigation concern computation and biology. Many computational models involve notions of memory, and one can examine these to find mechanical interpretations of these computational concepts. The most direct application of the mechanical formalism to simple computational models, such as finite automata, Turing machines, and random access machines, treats the entire machine state as the mass of the system, imputes a constant velocity to the system, and interprets all changes in terms of mass-flux forces that change the mass but not the velocity. The RMS models summarized in the preceding might point the way to computational models involving less trivial forms of mechanical concepts.

One might also examine biology for unappreciated mechanical structure. Molecular genetics, whether of the individual or of a population, provides one obvious starting point, with the genotype of an individual

constituting its biological mass and the phenotype constituting its biological position. Mutation, recombination, and reproduction possibly constitute actions in biological space. Admissible chemical processes of molecular constitution might then provide limits on the magnitudes of forces exerted in natural genetic processes.

Genetic algorithms interpret simple types of computational memories as genetic information that determines derived states. This recalls the relation of base reasons and derived states in the RMS. Mechanical interpretations of the RMS might thus prove readily adaptable to constructing mechanical interpretations of artificial genetic reproduction and variation, and so extend to interpreting actual genetic processes.

Experimental: Hybrid mechanics offers a mathematically coherent framework for understanding systems inhabiting different spaces but coupled through dynamogenesis. The existence of this mathematical framework says nothing about whether our world provides instances of such coupled mechanical systems. In particular, apart from the interaction of mind and body in individuals, we lack both demonstrable examples of physical forces produced by mental systems and proofs that such do not occur.

As mentioned earlier, Truesdell once published a brief query seeking evidence of measurable nonphysical forces on material objects [56]. The query in its entirety reads:

Can any reader supply examples of magic whose effect is measured? E.g., a magician whose spells could lift a ten-pound weight, but none heavier.

The query does not spell out whether Truesdell had in mind psychokinetic or supernatural forces, but it gives a model for the sort of evidence for which one might look. The difficulty, of course, lies in the expectation that such forces will stem from choices of conscious or free agents, who, like ornery people everywhere, need not follow strict regulations ensuring uniform experimental response.

Even if finding evidence for some possible types of hybrid mechanical systems in nature proves difficult, one can look for such systems in artificial systems as well. In particular, artificial intelligence provides examples in which these hybrid forces play crucial roles within the nonphysical factor spaces. The subfield of control theory dealing with hybrid systems provides models and real examples of systems moving in multiple disparate spaces, though these models do not give mechanical form to the nonphysical factor systems.

6.4 Precursors

Attempts to understand psychology and economics in terms of physical science have exploited many different mechanical analogies, include analogies to forces, energetics, thermodynamics, and machines. None of these attempts succeeded. One should not tar the present mechanical understanding of psychology and economics with the failures of these precursors, for the present understanding rests not on analogy but on reality, on showing how to understand these systems as actually having mass, exhibiting forces, and the like by virtue of satisfying the axioms defining those concepts.

From the perspective of the mechanical understanding of psychology and economics, one can say that earlier physically-inspired psychological theories failed for two main reasons. First, they lacked any precise

psychological theories, so the analogy bore the complete burden of providing a formalization. Without any independent verification of the psychological plausibility of the formalization so produced, the analogies failed to provide the desired understanding and prediction. Putting vague ideas into otherwise meaningless symbols benefits neither the vague ideas nor the pathetic symbols. Second, the attempts at physical analogy required no small suspension of belief due to the mismatch between the continuous models of physical behavior and the partially continuous but undeniably discrete nature of thought. No one disputes the importance of chemical dynamics and thermomechanics in baking a cake, but expressing the baker's recipe in differential equations—or even expecting such expression to convey understanding—is ludicrous.

Psychology: Herbart [24] made much of the “forces” conflicting concepts exert on each other, and devised a numerical scheme for calculating the magnitudes of these forces. He gave no formal mechanical basis within which to interpret these quantities—no bodies, no space, no motions—so the entire scheme moved only slightly away from mere suggestive analogy toward a quantitative theory of strength of belief. Subsequent psychologists dropped the mechanical conceptualization but modified the numerical schemes into theories of psychophysical measurement that formed the basis for the emerging experimental psychology: an outcome of no little irony, seeing how Herbart explicitly denied the very possibility of psychological experimentation.

A century later, Shand [50] also used an explicit mechanical terminology of forces and characteristic properties in discussing motivation, emotion, and character.

If we are to have a complete science of the mind, this will include a science of character as the most important part of it; and if we are to make any approach to such a science, it would seem that we must begin by a study of the fundamental emotions and of the instincts connected with them. But we have to conceive of the problem as essentially dynamical. The emotions are forces, and we have to study them as such. [50, p. 1]

We have then first to investigate the forces at the base of character, and the part they play in the general economy of mind.

The solution of this problem presupposes that we can profitably study the emotions dynamically, and that for this purpose we can sufficiently isolate them from one another and from the character as a whole. . . . In a strict sense we can never isolate the emotions. Each is bound up with others. Each subsists and works in a mental environment in which it is liable to be interfered with by the rest. Nor do these forces keep themselves, like human beings in the social environment, always distinct. On the contrary, they frequently become blended together, and often what we feel is a confused emotion which we cannot identify. [50, p. 2]

Shand provided no formalism comparable to standard mechanics or Herbart's quantification of forces, but presented detailed discussions of the roles different emotions and sentiments play in dynamogenesis. Each of these emotions generated characteristic dynamics, in his view, acting to increase the order exhibited in mental states.

Psychology produced numerous theories based on analogies to mechanical energy and motion determined by fields of potential energy. Freud [21] proposed an early theory involving energy flows between id,

ego, and superego, but this theory was always purely informal and suggestive, never respectable. Similar insubstantiality appears in Zipf's [65] work in linguistics, metaphorically connecting the structure of language to least-action principles in physics. Though this work proved premature as an application of physics to language, "Zipf's Law" relating word length and word frequency in language use represents an anticipation of important notions in information theory and statistics. Lewin's [30] topological theory of psychology cast mental dynamics in terms of potential field theory. Unfortunately, there was never any real substance to his theory: like Freud's, it relied mainly on purely suggestive terminology, but unlike Freud's, in a presentation laced with a few unsupported mathematical symbols to lend an air of mathematical rigor and meaning. The well-publicized catastrophe theory of Thom [55] and Zeeman [64] drew on respectable mathematical structures and represented talk by mathematicians who knew whereof they spoke mathematically, but conveyed only suggestive psychological concepts. With the mathematical structures providing no guidance to the applier, other than to choose the application so as to obtain the desired answer, catastrophe theory provided no more specifically psychological conclusions than did number theory (cf. [57]).

Economics: In economics, which stands on a much more extensive and fruitful mathematical basis than psychology, people have proposed both mechanical analogies and formal theories regularly for many years, to the point where Samuelson [46] complained of the flow of such "crank" papers onto his desk. It may well be that all were hopelessly flawed, justifying Samuelson's judgments of them, for the present theory provides the insight that the ideal rational agent studied by economists lacks the important mechanical property of inertia which underlies the limited abilities and resistance to change that constitute the most obvious differences between the imaginary species *homo economicus* and the actual species *homo sapiens sapiens*. Marshall [31] was justified in his mildly apologetic use of mechanical analogies, for he talked of people. Knight [29] was also justified in his criticism of Marshall's analogies, for Knight talked of an ideal.

Energetic analogies have largely proved vacuous due to the non-mathematical vacuity of the underlying variational theories, which rely only on abstract potential functions into which one can fit almost any theory of anything. Indeed, Feynman motivated variational theories in terms of an "unworldliness" function U , defined essentially to be the square of the difference between a conceivable system behavior and the "legal" behaviors of interest. These energetic theories lack even the suggestion of properly mechanical notions like mass and force. One encounters only potential fields and dynamical systems, general mathematical structures for describing almost any changing system as long as one can encode all the laws of the system, physical or otherwise, into the Lagrangian function (cf. [54]). This vacuity reduces these theories, when formal, to merely another way of writing systems of differential or difference equations. One can also apply the Hamiltonian formalism to economic dynamics (see [8]), but such applications merely recast economic equations as dynamical systems with conjugate dimensions (symplectic structures). Such conceptually-impooverished imitations of mechanical methods have done little to illuminate relations between mechanics and economics, and have brought disrepute upon the enterprise (see [34]).

Thermal theories and theories based on statistical mechanics may have not inspired any overall psychological theories, though they regularly figure in macroeconomic reporting, as in "economic competition heats up as the economy contracts and cools as it expands". Little formal application of these concepts

has appeared save application of methods of statistical mechanics to search problems where statistical techniques serve to introduce a measure of randomizing noise into the search, noise introduced to avoid the local minima that capture and stall standard gradient-following search methods. Search algorithms, however good, say little about mental organization and the structure of thinking. The behaviors of the algorithms instead derive their most important properties from those non-search characteristics of mind concerned with how one poses and designs the search problem.

Computation: The preceding analogies to chemistry, mechanical forces, energetics, and thermodynamics represent the failures, not the successes of physical analogy. The computational analogy to machines, based on purely kinematical concepts akin to clockwork gears, has delivered the major success of mechanical analogy to date.

The philosophers of the Enlightenment believed animals to be automata, merely complicated biological machines, but believed humans to have free souls controlling their bodies which, were the soul to depart, would exhibit only innate automatic behavior. The development of the Church-Turing theses of mechanizability of a wide range of procedures began to push back the boundary between mechanism and mind, to the point where Turing's [61] speculations on the mechanization of thought, capping two decades of cascading conceptual and technical advances by Turing, von Neumann, Wiener, McCulloch, and others, gave rise to the modern field of artificial intelligence, typically viewed as dating from the famed Dartmouth meeting of 1956.

The pallid kinematic picture based on automata theory that artificial intelligence employs in itself provides little help in explaining and understanding the structures of realistic psychologies or the nature of limited rationality. To gain more insight, we must look elsewhere, whether to the economics of computational resource allocation or to the full dynamical conception of mechanics.

7 Conclusion

There is no philosophy that is not founded upon knowledge of the phenomena, but to get any profit from this knowledge it is absolutely necessary to be a mathematician.

Daniel Bernoulli [5]

Mechanics remains a living subject, not a specimen frozen or embalmed for perpetual reference. Different areas of mechanics make do with different axioms, and our deformations of standard axioms, though themselves nontraditional, maintain this familiar variability of tradition. The mechanical formulation of psychology and economics opens many problems to new avenues of investigation, in these fields and in mathematics and mechanics as well.

Some historical perspective may help in facing up to the many unanswered questions raised by this work.

About three and a half centuries ago, three things happened. Descartes formulated a dualistic theory of mind, Newton and Leibniz invented the infinitesimal calculus, and Newton transformed mechanics from

a descriptive, philosophical subject to a mathematical subject based on an axiomatic perspective. Natural philosophy then made much progress for three hundred years, though it abandoned the initial form of each of these contributions. Philosophers discarded dualism for materialism and idealism because they thought mind-body interactions nonsensical, and idealism eventually gave way to materialist psychology and neurophysiological reductionism. Mathematicians abandoned infinitesimals because they seemed to obey no set rules, and Cauchy and Weierstrass showed how to develop analysis without them. Physicists abandoned Newton's axioms as vague and informal, especially after Euler invented what we now call Newton's equations to use in practical computations. Physicists even gradually abandoned mechanics as a subject of scientific study in favor of a reductionistic focus on field theories of smaller and smaller particles.

A funny thing then happened on the way to the future. About half a century ago, Walter Noll discovered how to axiomatize mechanics, and improved his initial axiomatization over the following decade or two into an axiomatic theory of great beauty and generality that revolutionized the theory and practice of mechanics. About the same time, Abraham Robinson realized that model theory provided means to make a proper mathematical theory of infinitesimals, and showed how to redevelop much of classical mathematics within nonstandard analysis, a result which then led to proofs of new results that had seemed too difficult to obtain in standard models of the real numbers. Also about the same time, Turing, von Neumann, and others began seriously investigating computational embodiments of thought that at least superficially exhibited the Cartesian dissimilarity from physical materials and which produced concrete models of rational thought precise enough to permit the present analysis in terms of slight modifications of Noll's axioms for mechanics.

It may well take time to see how to apply mechanical concepts to obtain improvements in scientific understanding and engineering power in psychology and economics. The centuries separating the original invention and recent formalization of ideas in mathematics, mechanics, and mind should give us motivation to continue these investigations even if we do not yet see clearly how or whether the present ideas will bear fruit. Hilbert [25] said it well, regardless of subsequent discoveries that showed his conviction requires minor modification unrelated to the present application:

This conviction of the solvability of every mathematical problem is a powerful incentive to the worker. We hear within us the perpetual call: There is the problem. Seek its solution. You can find it by pure reason, for in mathematics there is no *ignorabimus*.

Acknowledgment

This paper briefly summarizes ideas presented at length in *The Mechanical Foundations of Psychology and Economics*, a book first drafted in 1984 and currently nearing completion. I thank Carnegie Mellon University and the Massachusetts Institute of Technology for supporting me over the years of this writing, and thank Joseph Schatz, Gerald Sussman, Dana Scott, Richmond Thomason, Walter Noll, Peter Szolovits, Randall Davis, Robert Laddaga, and Joel Moses, for valuable advice and discussions. I especially thank Noll and the late Clifford Truesdell for papers and books that made this work possible.

References

- [1] P. Aczel. *Non-Well-Founded Sets*. Center for the Study of Language and Information, 1988.
- [2] R. Alur, C. Courcoubetis, T. A. Henzinger, and P.-H. Ho. Hybrid automata: an algorithmic approach to the specification and verification of hybrid systems. In R. Grossman, A. Nerode, A. Ravn, and H. Rischel, editors, *Hybrid Systems*, volume 736 of *Lecture Notes in Computer Science*, pages 209–229, Berlin, 1993. Springer-Verlag.
- [3] K. J. Arrow. *Social Choice and Individual Values*. Yale University Press, second edition, 1963.
- [4] N. D. Belnap. How a computer should think. In G. Ryle, editor, *Contemporary Aspects of Philosophy*, pages 30–56. Oriel Press, Stocksfield, 1976.
- [5] D. Bernoulli. Letter to John Bernoulli III. Dated 7 January 1763.
- [6] C. Boutilier, T. Dean, and S. Hanks. Decision theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 10:to appear, 1999.
- [7] M. S. Branicky. Universal computation and other capabilities of hybrid and continuous dynamical systems. *Theoretical Computer Science*, 138(1):67–100, Feb. 1995. Special issue on hybrid systems.
- [8] D. Cass and K. Shell. *The Hamiltonian Approach to Dynamic Economics*. Academic Press, New York, 1976.
- [9] J. M. Davoren. *Modal Logics for Continuous Dynamics*. PhD thesis, Cornell University, Ithaca, New York, January 1998.
- [10] J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12(2):231–272, 1979.
- [11] J. Doyle. A model for deliberation, action, and introspection. AI-TR 581, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 545 Technology Square, Cambridge, MA, 02139, 1980.
- [12] J. Doyle. Some theories of reasoned assumptions: An essay in rational psychology. Technical Report 83-125, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1983.
- [13] J. Doyle. The mechanical foundations of psychology and economics, with reflections on biology, computation, and artificial intelligence. Monograph in preparation, with major drafts in 1984, 1986, 1992, and 1999–2001, 1984.
- [14] J. Doyle. Expert systems and the “myth” of symbolic reasoning. *IEEE Transactions on Software Engineering*, SE-11(11):1386–1390, 1985.

- [15] J. Doyle. Artificial intelligence and rational self-government. Technical Report CS-88-124, Carnegie-Mellon University Computer Science Department, 1988.
- [16] J. Doyle. Constructive belief and rational representation. *Computational Intelligence*, 5(1):1–11, Feb. 1989.
- [17] J. Doyle. Rational belief revision. Presented at the Third International Workshop on Nonmonotonic Reasoning, Stanford Sierra Camp, CA, available at <http://www.medg.lcs.mit.edu/doyle/ftp/rbr90.ps>, June 1990.
- [18] J. Doyle. Reasoned assumptions and rational psychology. *Fundamenta Informaticae*, 20(1-3):35–73, 1994.
- [19] J. Doyle and M. P. Wellman. Rational distributed reason maintenance for planning and replanning of large-scale activities. In K. Sycara, editor, *Proceedings of the DARPA Workshop on Planning and Scheduling*, pages 28–36, San Mateo, CA, Nov. 1990. Morgan Kaufmann.
- [20] J. Doyle and M. P. Wellman. Impediments to universal preference-based default theories. *Artificial Intelligence*, 49(1-3):97–128, May 1991.
- [21] S. Freud. Project for a scientific psychology. In J. Strachey, editor, *Standard Edition of the Complete Psychological Works of Sigmund Freud*, volume I. Hogarth Press, London, 1895. Reprinted in 1966.
- [22] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In M. Y. Vardi, editor, *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 83–95, Los Altos, CA, March 1988. Morgan Kaufmann.
- [23] J. Hartmanis. Turing award lecture on computational complexity and the nature of computer science. *Commun. ACM*, 37(10):37–43, oct 1994.
- [24] J. F. Herbart. *A Text-Book in Psychology: An Attempt to Found the Science of Psychology on Experience, Metaphysics, and Mathematics*. D. Appleton, New York, 1891. Translation of the second revised edition of 1834 (Hamburg and Leipzig, G. Hartenstein, 1886). First edition 1816.
- [25] D. Hilbert. Mathematical problems. *Bulletin of the American Mathematical Society*, 8:437–479, 1902. Lecture delivered before the 1900 International Congress of Mathematicians at Paris.
- [26] W. James. *The Principles of Psychology*. Henry Holt & Co., New York, 1890. In two volumes.
- [27] W. James. *The Will to Believe and Other Essays in Popular Philosophy*. Longmans, Green, and Co., New York, 1897.
- [28] R. E. Kalman, P. L. Falb, and M. A. Arbib. *Topics in Mathematical System Theory*. McGraw-Hill, New York, 1969.

- [29] F. H. Knight. Statics and dynamics: some queries regarding the mechanical analogy in economics. In *On the History and Method of Economics*, pages 179–201. University of Chicago Press, Chicago, 1956.
- [30] K. Lewin. *Field Theory in Social Science*. Harper, New York, 1951.
- [31] A. Marshall. *Principles of Economics*. Macmillan, New York, eighth edition, 1949.
- [32] J. C. C. McKinsey, A. C. Sugar, and P. Suppes. Axiomatic foundations of classical particle mechanics. *Journal of Rational Mechanics and Analysis*, 2:253–272, 1953.
- [33] M. Minsky. *The Society of Mind*. Simon and Schuster, New York, 1986.
- [34] P. Mirowski. *More Heat Than Light: Economics as social physics: Physics as nature's economics*. Cambridge University Press, Cambridge, 1989.
- [35] W. Noll. The foundations of classical mechanics in the light of recent advances in continuum mechanics. In P. S. L. Henkin and A. Tarski, editors, *The axiomatic method, with special reference to geometry and physics. Proceedings of an International Symposium held at the University of California, Berkeley, December 26, 1957–January 4, 1958*, Studies in Logic and the Foundations of Mathematics, pages 266–281, Amsterdam, 1958. North-Holland. Reprinted in [39].
- [36] W. Noll. La mécanique classique, basée sur un axiome d'objectivité. In *La Méthode Axiomatique dans les Mécaniques Classiques et Nouvelles (Actes Quatrième Colloq. Internat. de Logique et Philos. Sci., Paris, 1959)*, pages 47–56, Paris, 1963. Gauthier-Villars. Reprinted in [39].
- [37] W. Noll. A new mathematical theory of simple materials. *Archive for Rational Mechanics and Analysis*, 48:1–50, 1972. Reprinted in [39].
- [38] W. Noll. Lectures on the foundations of continuum mechanics and thermodynamics. *Archive for Rational Mechanics and Analysis*, 52:62–92, 1973. Reprinted in [39].
- [39] W. Noll. *The Foundations of Mechanics and Thermodynamics: Selected Papers*. Springer-Verlag, Berlin, 1974.
- [40] W. Noll. On material frame-indifference. *Archive for Rational Mechanics and Analysis*, 1996. to appear.
- [41] W. Noll and E. G. Virga. Fit regions and functions of bounded variation. *Archive for Rational Mechanics and Analysis*, 102:1–21, 1988.
- [42] B. Pascal. *Pensées sur la religion et sur quelques autres sujets*. Harvill, London, 1962. Translated by M. Turnell, originally published 1662.

- [43] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [44] R. Penrose. *The Emperor's New Mind: Concerning Computers, Minds, and The Laws of Physics*. Oxford University Press, New York, 1989.
- [45] B. Russell. *The Conquest of Happiness*. Liveright, New York, 1930.
- [46] P. A. Samuelson. Maximum principles in analytical economics. *Science*, 173(4001):991–997, September 10 1971.
- [47] L. J. Savage. *The Foundations of Statistics*. Dover Publications, New York, second edition, 1972.
- [48] D. Scott. Models for various type-free calculi. In P. Suppes, L. Henkin, A. Joja, and G. C. Moisil, editors, *Logic, Methodology and Philosophy of Science IV*, pages 157–187. North-Holland, Amsterdam, 1973.
- [49] D. S. Scott. Domains for denotational semantics. In M. Nielsen and E. M. Schmidt, editors, *Automata, Languages, and Programming: Ninth Colloquium*, volume 140 of *Lecture Notes in Computer Science*, pages 577–613, Berlin, 1982. Springer-Verlag.
- [50] A. F. Shand. *The Foundations of Character: Being a Study of the Tendencies of the Emotions and Sentiments*. Macmillan and Company, London, second edition, 1920.
- [51] C. E. Shannon. A mathematical theory of communication. *Bell System Tech. J.*, 27:379–423, 623–656, 1948.
- [52] H. A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118, 1955.
- [53] H. A. Simon. Rational choice and the structure of the environment. *Psychological Review*, 63:129–138, 1956.
- [54] G. J. Sussman, J. Wisdom, and M. E. Mayer. *Structure and Interpretation of Classical Mechanics*. MIT Press, Cambridge, Massachusetts, 2001.
- [55] R. Thom. *Structural Stability and Morphogenesis: An outline of a general theory of models*. Benjamin, Reading, 1975.
- [56] C. Truesdell. Bounded magic. *Isis*, 47:59, 1956. Query No. 150.
- [57] C. Truesdell. Our debt to the French tradition: “catastrophes” and our search for structure today. In *An Idiot's Fugitive Essays on Science: Methods, Criticism, Training, Circumstances*, pages 80–94. Springer-Verlag, New York, 1984.

- [58] C. Truesdell. Suppesian stews. In *An Idiot's Fugitive Essays on Science: Methods, Criticism, Training, Circumstances*. Springer-Verlag, New York, 1984.
- [59] C. Truesdell and W. Noll. *The Non-Linear Field Theories of Mechanics*. Springer-Verlag, Berlin, second edition, 1992.
- [60] C. A. Truesdell, III. *A First Course in Rational Continuum Mechanics*, volume 1: General Concepts. Academic Press, New York, second edition, 1991.
- [61] A. M. Turing. Computing machinery and intelligence. *Mind*, 49:433–460, 1950.
- [62] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, third edition, 1953.
- [63] B. T. Yates. *Self-Management: The Science and Art of Helping Yourself*. Wadsworth, Belmont, CA, 1985.
- [64] E. C. Zeeman. *Catastrophe Theory: Selected Papers 1972-1977*. Addison-Wesley, London, 1977.
- [65] G. K. Zipf. *Human Behavior and the Principle of Least Effort: An introduction to human ecology*. Addison-Wesley, Cambridge, 1949.