Are Video Games Good for Learning?

Keynote address at Curriculum Corporation 13th National Conference, Adelaide, August 2006

James Paul Gee
Tashia Morgridge Professor of Reading
University of Wisconsin-Madison, United States

jgee@education.wisc.edu

An emerging research field relates to the hypothesis that video games are good for learning (Shaffer, Squire, Halverson & Gee in press; Gee 2003, 2005). This hypothesis amounts to two claims. The first is that good commercial games are built on sound learning principles (Gee 2003) that are supported by research in the learning sciences (Bransford, Brown and Cocking 2000). The second is that video game technologies hold out great promise for moving beyond entertainment, to building new learning systems for serious purposes in and out of school.

A rough distinction can be made between two major types of games: what I call ‘problem games’ and ‘world games’. Problem games, such as *Tetris* and *Diner Dash*, focus on solving a given problem or class of problems, while world games, such as *Half-Life* and *Rise of Nations*, simulate a wider world within which the player may solve many different sorts of problems. Classic early arcade games like *Super Mario* sit right at the border of this division. Among the different types of video games, some are likely to be better or worse suited to meeting different learning goals.

Features of games with high learning potential

What are the features of video games that hold out promise for good learning? I will sketch some of these features and also suggest accompanying research questions. I start with features that relate equally well to non-game experiences, and then turn to features that are more directly related to the ‘gameness’ of video games. However, we should keep in mind that the ‘non-game’ features may *not* work as well for learning if they are detached from the ‘game’ features. That is *one* key question for research.

1. Empathy for a complex system

Consider scientific simulations of phenomena such as weather systems, atoms, cells or the rise and fall of civilisations. Scientists are not ‘inside’ these simulations in the way that players are ‘inside’ the simulated worlds of games like *Thief*. The scientist doesn’t ‘play’ an ant in his or her simulation of an eco-system. The scientist doesn’t discover and form goals from the perspective of the ant in the way that a player takes on the perspective of Garrett in *Thief*.

However, at the cutting edge of science, scientists often talk and think *as if* they were ‘inside’ not only the simulations they build but even the graphs they draw. They try to think from ‘within’ local regions of the system being simulated, while still keeping in mind the system as a whole. In this way they aim to gain a deeper feel for how variables are interacting within the system. Just as a player ‘becomes’ Garrett, a scientist may talk and think as if he or she actually *were* an electron in a certain state or an ant in a colony. For example, consider the following well-known example of a physicist talking to other physicists while looking and pointing to a graph on a blackboard:

> But as you go below the first order transition you’re (*leans upper body to right*) still in the domain structure and you’re still trying to get (*sweeps right arm to left*) out of it. Well you also said (*moves to board; points to*

> *diagram*) the same thing must happen here (*points to the right side of the diagram*). When (*moves finger to left*) I come down (*moves finger to right*) I'm in (*moves finger to left*) the domain state (Ochs, Gonzales & Jacoby 1996, pp. 330–1).

Notice the use of 'you' and 'I'. The scientist talks and acts as if he and his colleagues are moving their bodies not only inside the graph, but also inside the complex system it represents.

A key research question that arises here is this: though video games and scientific simulations are not the same thing, can video games, under the right circumstances, encourage and actually enact a similar 'attitude' or 'stance' to that taken by scientists who study complex systems? This stance involves a sort of embodied empathy for a complex system, where a person seeks to participate in and within a system, all the while seeing and thinking of it as a system rather than just a set of local or random events. This *does* seem to be the sort of stance players take when playing the role of Garrett in a game like *Thief* in that they seek to figure out the rule system that underlies the virtual world through which Garrett – and they – move. Could video games create this kind of empathy for the sorts of complex systems relevant to academic and other domains that lie beyond entertainment, such as urban planning, space exploration or global peace?

2. Simulations of experience and preparations for action

Video games do not just carry the potential to replicate a sophisticated scientific way of thinking: they may actually externalise, in a better fashion than any other technology we currently have, the ways in which the human mind works and thinks. Consider, in this regard, some recent research in the learning sciences:

> … comprehension [understanding words, actions, events, or things] is grounded in perceptual simulations that prepare agents for situated action (Barsalou 1999a, p. 77)

> … higher intelligence is not a different kind of process from perceptual intelligence (Hawkins 2004, p. 96).

In these remarks, human understanding is not viewed primarily as a matter of storing general concepts in the head or applying abstract rules to experience. Rather, the view is that humans think and understand best when they can imagine (simulate) an experience in such a way that the simulation prepares them for actions they need and want to take in order to accomplish their goals (Barsalou 1999b; Clark 1997; Glenberg & Robertson 1999). Effective thinking is viewed as the human actor seeing how the world, at a specific time and place (as it is given, but modifiable), can afford the opportunity for actions that will lead to a successful accomplishment of the actor's goals. Any generalisations are formed on the basis of experience and imagination of experience.

Video games are external (not theoretical) simulations of worlds or problem spaces in which the player. using a particular perspective, must prepare for action and the accomplishment of goals. Gamers learn to see the world of each different game in a way that is distinctive to that game; but every game requires its players to see the virtual world in terms of how it will afford the sorts of actions they ('they' meaning a melding of themselves and their virtual character) need to take to accomplish their goals (to win in the short and long run).

For example, players see the world in *Full Spectrum Warrior* as routes (for a squad) between cover (for example, corner to corner, house to house), because this prepares them for the actions they need to take, namely attacking without being vulnerable to attack. They see the world of *Thief* in terms of light and dark, illumination and shadows, because this prepares them for the different actions they need to take in this world, namely hiding, disappearing into the shadows, sneaking and in various other ways moving unseen to their goal.

If commercial video games often offer worlds in which players prepare for the actions of soldiers or thieves, could other types of games let players prepare for action from different perspectives or identities, for example, a particular type of scientist, political activist or global citizen? If games could do so, they would speak to one of the deeper problems of education: that many students who can pass paper-and-pencil tests cannot apply their knowledge to solving real problems (Gardner 1991).

3. Distributed intelligence via the creation of smart tools

Good video games distribute intelligence (Brown, Collins, & Dugid 1989) between a real-world person and artificially intelligent virtual characters. For example, in *Full Spectrum Warrior*, the player uses the buttons on the controller to give orders to two squads of soldiers (the game *SWAT 4* is a similar example). The instruction manual for the game makes it clear at the outset that, to play the game successfully, players must take on the values, identities and ways of thinking of a professional soldier: 'Everything about your squad,' the manual explains, 'is the result of careful planning and years of experience on the battlefield. Respect that experience, soldier, since it's what will keep your soldiers alive' (p. 2). In the game, that experience – the skills and knowledge of professional military expertise – is distributed between the virtual soldiers and the real-world player. The soldiers in the player's squads have been trained in movement formations; the role of the player is to select the best position for them on the field. The virtual characters (the soldiers) know part of the task (various movement formations) and the player must come to know another part (when and where to engage in such formations). This kind of distribution holds for every aspect of professional military knowledge in the game.

By distributing knowledge and skills this way – between the virtual characters (smart tools) and the real-world player – the player is guided and supported by the knowledge built into the virtual soldiers. This offloads some of the cognitive burden from the learner to smart tools that can do more than the learner is currently capable of doing by himself or herself. It allows the player to begin to act, with some degree of effectiveness, before being really competent: 'performance before competence'.

The player gains competence through trial, error and feedback, rather than having to wade through a lot of text before being able to engage in activity.

Such distribution also allows players to internalise not only the knowledge and skills of a professional (a professional soldier in this case), but also the concomitant values ('doctrine' as the military says) that shape and explain how and why that knowledge is developed and applied in the world. This suggests an important question for research: could other occupational roles, for example, scientists, doctors, government officials, urban planners and political activists (Shaffer 2004), be modeled and distributed in this fashion as a deep form of value-laden learning? Could learners come to compare and contrast different value systems as they play different games?

Shaffer's (2004; 2005) 'epistemic games' already give us a good indication that even young learners, through video games embedded inside a well-organised curriculum, can be inducted into professional practices as a form of value-laden deep learning that transfers to school-based skills and conceptual understandings. However, much work remains to be done in making the case that the knowledge, skills and values that may be developed within such games are transferable to other areas of school learning and, in particular, to students' learning in traditional content areas.

4. Cross-functional teamwork

Consider a small group partying (hunting and questing) together in a massive multiplayer game like *World of WarCraft*. The group might well be composed of a hunter, warrior, druid, mage, and priest. Each of these character types has quite different skills and plays the game in a different way. Each group member must learn to be good at his or her special skills and, as a team member, to integrate these skills within the group as a whole. In order to achieve a successful integration, each team member must share with all the other members of the group some common knowledge about the game and game play, including some understanding of the specialist skills of other player types. Each member of the group must have specialist knowledge (intensive knowledge) and also general common knowledge (extensive knowledge), including knowledge of the other member's functions.

Players – who are interacting with each other in the game and via a chat system – orient to each other not in terms of their real-world race, class, culture or gender (these may very well be unknown or, if communicated, be fictional) but, first and foremost, through their identities as game players and, in particular, as players of *World of WarCraft*. They can also use their real-world race, class, culture and gender (for better or worse) as strategic resources if and when they please, and the group can draw on the differential real-world resources of each player, but in ways that do not force them as players into pre-set racial, gender, cultural or class categories.

It has been argued that this form of affiliation – what I call cross-functional affiliation – is crucial for the workplace teams of modern 'new capitalist' workplaces, and also for contemporary forms of social activism such as that of the Green movement (Beck 1999; Gee 2004; Gee, Hull & Lankshear 1996). People specialise, but also integrate and share. They organise around a primary affiliation to their common goals and use their cultural and social differences as strategic resources, not as barriers. The crucial research questions that arise are the extent to which such collaborative work in commercial games could transfer to collaborative abilities in other settings; and the extent to which good video games designed with different content can teach collaboration and cross-functional teamwork for institutions like schools and workplaces. The United States Army has made this assumption in regard to games like *America's Army*, but it remains to be empirically demonstrated in other domains.

5. Situated meaning

Words have more than just general dictionary-like meanings. They can have different and specific meanings in the particular situations in which they are used and in the different specialist domains that recruit them (Gee 2004). This is true of the most mundane cases. For instance, notice the change in meaning in the word 'coffee' when used in the following different situations:

'The coffee spilled, go get the mop' (coffee as liquid).

'The coffee spilled, go get a broom' (coffee as grains).

'The coffee spilled, stack it again' (coffee in cans).

Or notice the quite different meanings of the word 'work' when used in everyday life as opposed to its use in physics. For example, in everyday life, I can say that I worked hard to push the car; but if my efforts did not move the car, I did *no* 'work' in the sense in which that word is used in physics.

A good deal of school success is based on being able to understand complex academic language (Gee 2004), as illustrated in this example from a high-school science textbook.

> The destruction of a land surface by the combined effects of abrasion and removal of weathered material by transporting agents is called erosion ... The production of rock waste by mechanical processes and chemical changes is called weathering.

When students understand such language only verbally, rather than in a situated fashion, they can trade words for words, replacing the words with their definitions. However, while they may be able to pass paper-and-pencil tests, they often can't use the complex language of the text for problem solving, because they don't actually understand how the language applies to specific cases. We have known for years now that a great many school students can get good grades on paper-and-pencil tests in science, for example, but can't use their knowledge to solve actual problems (Gardner 1991). Students come to understand the words in a situated fashion only if and when they can apply the words to specific situations and to the solution of specific problems.

People acquire situated meanings for words – that is, meanings that they can apply in actual contexts of use for action and problem solving – only when they have heard these words in interactional dialogue with people more expert than themselves (Tomasello 1999) and when they have experienced the images and actions to which the words apply (Gee 2004). Dialogue, experience and action are crucial if people are to have more than just words for words, if they are to be able to relate words to actual experiences, actions, functions and problem solving. As they develop this ability within an increasing number of contexts, they are able to increasingly generalise the meanings of the word, but the words never lose their moorings in talk, embodied experience, action and problem solving.

Since video games are simulations of experience, they can put language into the context of dialogue, experience, images and actions. They allow language to be situated. Furthermore, good video games give verbal information 'just in time' – near the time it can actually be used – or 'on demand', when the player feels a need for it and is ready for it (Gee 2003). They do not give players lots and lots of words out of context before they can be used and experienced or before they are needed or useful. Video games provide players with a context and a need for acquiring new words and new forms of language for new types of activity, whether players be members of a SWAT team or scientists. Given the importance of oral and written language development (for example, vocabulary) to school success, it is crucial that this assumption be tested in terms of the language players pick up from commercial games (consider, for example, how young children play *Yu-Gi-Oh*, a game that contains very complex language indeed) and also in terms of how games can be made and used for the development of specifically school-based (or other institutional) language demands.

6. Open-endedness: melding the personal and the social

In a video game, the player 'plays' a character or set of characters. The player must discover the goals of this character within the game world and pursue those goals, using whatever abilities the character possesses. In *Thief*, the player comes to realise that reaching Garrett's goals requires stealth (for which Garrett is well suited). These are the 'in game' goals the player must discover and pursue.

But in good open-ended games, such as *The Elder Scrolls III: Morrowind*, *Arcanum*, *The Sims*, *Deus Ex 2*, *Mercenaries*, *Grand Theft Auto* and many others, players also construct their own goals, which are based on their own desires, styles and backgrounds. The player then attributes these personal goals to the virtual character and considers the affordances in the virtual world (by figuring out the rule system) for realising these personal goals, along with the virtual character's 'in game' goals.

For example, in *The Elder Scrolls III: Morrowind*, a player may decide to eschew heavy armour and lots of fighting in favour of persuasive skills, stealth and magic; or the player can engage in lots of face-to-face combat in heavy armour. The player can carry out a linear sequence of quests set by the game's designers or can make up his or her own quests, becoming so powerful that the designer's quests become easy and only a background feature of the game. In *Grand Theft Auto III*, the player can be evil or good (for example, by jumping into ambulances and doing good deeds), can undertake quests in different sequences, and can choose whether or not to play large some major sections of the game (for example, a player can trigger gang wars or avoid them altogether). Even in less open-ended games, even quite young players set their own standards of accomplishment, replaying parts of the game so that their hero pulls things off in the heroic fashion and style the player deems appropriate. This marriage of personal goals and ‘in game’ goals produces in players a state of high motivation.

When people are learning or doing science, they must discover and realise goals that are set up by the scientific enterprise as a domain and as a social community. These are equivalent to ‘in game’ goals. Effective learners, however, are those who marry these goals to their own personal goals, which are based on their own desires, styles and backgrounds, to arrive at some seamlessness between their scientific identity and their everyday world of personal and community-based identities and values. Just as good video games readily allow such a marriage, so should good science instruction. The research question here is: how could we use video games to achieve a marriage of ‘in game’ goals (the goals that flow from an academic area or from the teacher) with students’ personal goals and learning styles, for use in school learning and for learning in other contexts?

The six features: summary

So far I have discussed six features of good video games that I believe facilitate good learning; but they are features that are not particular to video games. These features are, to summarise:

1. Video games can create an embodied empathy for a complex system.
2. They are simulations of embodied experience and preparations for action.
3. They involve distributed intelligence via the creation of smart tools.
4. They create opportunities for cross-functional affiliation.
5. They allow meanings to be situated.
6. They can be open-ended in ways that encourage a melding of personal and social goals.

None of this is to say that video games do these good things all by themselves. It all depends on how they are used and what sorts of wider learning systems (activities and relationships) they are embedded within. And this, indeed, raises one of the most important research questions for the field of games and learning: what sorts of wider learning systems ought games to be embedded within if we are to leverage their powers for learning to the greatest extent? With what other activities – ‘in game’ and ‘out of game’ – ought they to be paired? What are the most effective roles for teachers in these learning systems?

Features of a good game

The features examined so far all seem, however, to be somewhat removed from the pleasures we derive from playing these games. An important question for research is how to embed features in video games that work well as games. What are the features that make a video game a good game? What are the sources of the pleasure we draw from video games? How can these features relate to effective learning?

1. Motivation

It is clear to see how profoundly motivating video games are for players. Players focus intently on game play for hours at a time, solving complex problems all along the way. In an 'attentional economy', where diverse products and messages, not to mention school subjects, compete for people's limited attention, video games can draw our deep attention. It is important that research be directed to understanding the source or sources of this motivation, if it is to provide a foundation for learning.

One hypothesis in this regard is that problem solving and learning, along with display of mastery, are themselves a key source of motivation in good video games. If this is true, then, why is learning and mastery so motivating in *this* context and not always as motivating within school? A second hypothesis is that learning and mastery are motivating in good video games because they make use of deep learning principles, including the six features we have just discussed. If so, we need to know whether the same level of motivation can occur when the content of a video game is based on more academic or specialised learning goals. Some people assume that science, for example, can never be made as enticing as fighting fictional wars in *America's Army* or running a family in *The Sims* – but I know of no good scientist who does not find science motivating, entertaining and life enhancing.

2. The role of failure

There are certainly features of video games *as games* that help explain the motivation they recruit and the learning they enable. One of these is the role of failure. In good games, the price of failure is lowered. When players fail, they can, for example, start again at their last saved game. Furthermore, failure – for example, failure to kill a boss – is often seen as a way to learn the underlying pattern and eventually to win. These features of failure in games allow players to take risks and try out hypotheses that might be too costly in places such as classrooms where the cost of failure is higher or where no learning stems from failure.

3. Competition and collaboration

Every gamer and game scholar knows that a great many gamers, including young ones, enjoy competition with other players in games, either one-on-one or team-based. It is striking that many young gamers see competition as pleasurable and motivating in video games, but not in school. Why this is so ought to be a leading question for research into games and learning. What seems evident is that competition in video games is seen by gamers as social and is often organised in ways that allow people to compete with people at their own level or as part and parcel of a social relationship that is as much about gaming as winning and losing. Furthermore, gamers highly value collaborative play, as in, for example, two people playing *Halo* together to beat the game or the group collaboration in massive multiplayer games like *World of WarCraft*. Indeed, collaboration and competition seem often to be closely related and integrated in gaming, whereas that is not usually the case in school.

4. The design of games

Beyond issues of motivation, failure, competition and collaboration, the very ways in which games are designed as games seem to give them features that promote learning and a sense of mastery. While this is a hypothesis that needs testing, some features of the design of video games appear to be closely associated with well-known principles of learning. Consider the following seven such design features:

(a) Interactivity

In a video game, players make things happen; they don't just consume what the 'author' (game designer) has placed before them. In good games, players feel that their actions and decisions – and not just the designers' actions and decisions – are co-creating the world of the game and the experiences they are having. Each player's actions matter and each player – based on his or her own style, decisions and actions – takes a somewhat different trajectory through the game world. All players engage in a form of simultaneous 'reading' (interpreting) and 'writing' (producing). The more open-ended the game, the more this is the case. All deep learning involves learners feeling a strong sense of ownership and agency, as well as the ability to produce and not just passively consume knowledge.

(b) Customisation

In some games, players are able to customise the game play to fit their learning and playing styles, for example through adopting different difficulty levels or choosing to play different characters with different skills. Other games are designed to allow a variety of styles of learning and playing, by providing multiple ways to solve the problems in the game, as, for example, in the *Deus Ex* games and in role-playing games like *Arcanum*. Customisation, in the sense of catering for a variety of learning styles and providing multiple routes to success, is an important learning principle in many contexts.

(c) Strong identities

Good games offer players identities that trigger a deep investment on the part of the player. Such identities are often connected to a specific virtual character, though sometimes to a whole 'civilisation' (as in *Civilisation* or *Rise of Nations*). When gamers are playing characters, strong identities are achieved through the character being so intriguing that players want to inhabit the character and can readily project their own fantasies, desires and pleasures onto the character (for example, Solid Snake in the *Metal Gear Solid* games); or through the player having to determine the traits of a relatively empty character in such a way that the player can create a deep and consequential life history in the game world for the character (for example, in role-playing games like *The Elder Scrolls III: Morrowind*). Furthermore, in games, the identity of the character one plays is very clearly associated with the sorts of functions, skills and goals one has to carry out in the virtual world. Many people have argued that this sense of identity (for example, 'being–doing a scientist' in order to learn science) is crucial for deep learning (see, for example, Gee 2004; diSessa 2000; Shaffer 2004).

(d) Well-sequenced problems

In good games the presentation of problems is carefully sequenced. In particular, some problems are introduced early in the game so as to lead players to form good guesses about how to proceed when they face harder problems later on in the game. In this sense, earlier parts of a good game are always looking forward to later parts. In connectionist approaches to learning, it is argued that such sequencing is crucial for effective learning in complex domains (see, for example, Elman 1991a, 1991b).

(e) A pleasant level of frustration

Good games adjust challenges and give feedback in such a way that a range of players can experience the game as challenging but doable and feel that their effort is paying off. Players get feedback that indicates whether or not they are on the right road for success later on and at the end of the game. When players lose to a boss, perhaps multiple times, they get feedback about the sort of progress they are making so that at least they know if and how they are moving in the right direction towards success. DiSessa (2000) has argued that such pleasant frustration is an optimal state for learning in areas such as science.

(f) A cycle of expertise

Good games create and support what has been called in the learning sciences the ‘cycle of expertise’ (Bereiter & Scardamalia 1993), with repeated cycles of extended practice and tests of mastery of that practice before a new challenge leads to new practice and new mastery. This is part of what constitutes good pacing in a game.

(g) ‘Deep’ and ‘fair’

These terms are used in the gaming community to describe the art of game construction. A game is ‘fair’ when it is challenging, but set up in a way that leads to success rather than building in failure over which the player has little or no control. A game is ‘deep’ when game play elements (for example, a fighting system in a turn-taking game) that initially seem simple, and so easy to learn and use, become more and more complex the more the player comes to master and understand them. These two characteristics might also be put to good use in the learning sciences, too.

Conclusion

These basic features of *games as games* appear to be important features of effective learning. It therefore seems worth exploring how they can be used, in association with the set of six features we discussed earlier as having high potential for learning – the ones less particular to games *as games* – to deliver new and highly effective learning systems for serious purposes in and out of school.

References

Barsalou, LW 1999a, 'Language comprehension: Archival memory or preparation for situated action', *Discourse Processes*, vol. 28, pp. 61–80.

Barsalou, LW 1999b, 'Perceptual symbol systems', *Behavioral and Brain Sciences*, vol. 22, pp. 577-660.

Bereiter, C & Scardamalia, M 1993, *Surpassing ourselves: An inquiry into the nature and implications of expertise*, Open Court, Chicago.

Beck, U 1999, *World risk society*, Blackwell, Oxford, UK.

Bransford, J, Brown, AL & Cocking, R 2000, *How people learn: Brain, mind, experience, and school*, expanded edn, National Academy Press, Washington, DC.

Brown, JS, Collins, A & Dugid1989, 'Situated cognition and the culture of learning', *Educational Researcher*, vol. 18, pp. 32–42.

Clark, A 1997, *Being there: Putting brain, body and world together again*, MIT Press, Cambridge, Mass.

diSessa, AA 2000, *Changing minds: Computers, learning and literacy*, MIT Press, Cambridge, Mass.

Elman, J 1991a, 'Distributed representations, simple recurrent networks and grammatical structure', *Machine Learning*, vol. 7, pp. 195–225.

Elman, J 1991b, 'Incremental learning, or the importance of starting small', Technical Report 9101, Center for Research in Language, University of California at San Diego.

Gardner, H 1991, *The unschooled mind: How children think and how schools should teach*, Basic Books, New York.

Gee, JP 2003, *What video games have to teach us about learning and literacy*, Palgrave/Macmillan, New York.

Gee, JP 2004, *Situated language and learning: A critique of traditional schooling*, Routledge, London.

Gee, JP 2005, *Why video games are good for your soul: Pleasure and learning*, Common Ground, Melbourne.

Gee, JP, Hull, G & Lankshear, C 1996, *The new work order: Behind the language of the new capitalism*, Westview, Boulder, CO.

Glenberg, AM & Robertson, DA 1999, 'Indexical understanding of instructions', *Discourse Processes*, vol. 28, pp. 1–26.

Hawkins, J (with S Blakeslee) 2004, *Intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machine*s, Times Books, New York.

Ochs, E, Gonzales, P & Jacoby, S 1996, 'When I come down I'm in the domain state', in E Ochs, E Schegloff & SA Thompson (eds), *Interaction and grammar*, Cambridge University Press, Cambridge, pp. 328–69.

Shaffer, DW 2004, 'Pedagogical praxis: The professions as models for post-industrial education', *Teachers College Record*, vol. 10, pp. 1401–21.

Shaffer, DW 2005, 'Epistemic games', *Innovate*, vol. 1, no. 6, http://www.innovateonline.info/index.php?view=article&id=81.

Shaffer, DW, Squire, K, Halverson, R & Gee, JP (in press), 'Video games and the future of learning', *Phi Delta Kappan*.

Tomasello, M 1999. *The cultural origins of human cognition*, Harvard University Press, Cambridge, MA.