11

NEUROSCIENTIFIC APPROACHES TO SELF-REGULATORY CONTROL IN SPORTS

Wanja Wolff, Anna Hirsch, Maik Bieleke, and Amitai Shenhav

Introduction

Sports performance critically depends on physical fitness and skill level. However, success also hinges on how well athletes deal with psychological obstacles that threaten optimal performance (Wolff, Bieleke, & Schüler, 2019). For example, in order to make a critical free-throw, basketball players have to ward off challenges to effective movement execution, triggered both externally (e.g. distraction by hostile chants) and internally (e.g. fear of failure). In the same vein, cyclists whose muscles are aching need to overcome their automatic inclination to ease up and slow down. These exemplary challenges are well within the scope of current definitions of self-regulatory control[1] as the "efforts people exert to stimulate desirable responses and inhibit undesirable responses" (de Ridder, Lensvelt-Mulders, Finkenauer, Stok, & Baumeister, 2012, p. 77) or "the set of mechanisms required to pursue a goal, especially when distraction and/or strong (e.g., habitual) competing responses must be overcome" (Shenhav, Botvinick, & Cohen, 2013, p. 217). Thus, it seems plausible that, in addition to fitness and skill, sports performance hinges on the successful exertion of mental effort in the service of self-regulatory control.

Indeed, a large body of research has found support for the importance of self-regulatory control in sports (Englert, 2016), both at the trait level (i.e. stable self-regulatory tendencies within an individual) and at the state level (i.e. situationally defined influences on one's tendency to self-regulate). For instance, on the trait level, one study has shown that elite cyclists outperform non-elite peers in a reaction-time-based measure of trait self-control (Martin et al., 2016). In another study, youth athletes in an elite talent development program scored higher on a self-report measure of self-regulatory control than their non-selected peers (Wolff, Bertrams, & Schüler, 2019). In the recreational context, high self-regulatory control has been linked with better adherence to exercise regimens (Stork, Graham, Bray, & Martin Ginis, 2016). This makes intuitive sense, as self-regulatory control has been shown to help people deal with the difficulties of adhering to valued behaviors (Wolff, Martarelli, Schüler, & Bieleke, 2020). At the state level, the application of self-regulatory control detrimentally affects subsequent sports performance (for a meta-analysis, please see Giboin & Wolff, 2019, but see also Holgado, Troya, Perales, Vadillo, & Sanabria, 2019). For example, impaired performance after prior self-regulatory control has been found in such diverse sport settings as sprint-running (Englert, Persaud, Oudejans, & Bertrams, 2015) and dart throwing (Yang, Park, & Shin, 2019). Finally, it has been shown that physical effort causes feelings of mental exertion, impairs self-regulatory control, and leads to hypoactivations in control-relevant areas of the brain (Blain et al., 2019; Wolff, Schüler et al., 2019). This is consistent with the claim that physical performance requires self-regulatory control.

A multitude of theoretical accounts have been proposed that specify *why* and *when* self-regulatory control is applied and why it sometimes appears to *fail* (Beedie & Lane, 2012; Inzlicht, Schmeichel, & Macrae, 2014; Kotabe & Hofmann, 2015; Kurzban, Duckworth, Kable, & Myers, 2013; Shenhav et al., 2013). In particular, recent years have seen substantial advancements in our understanding of the guiding principles of self-regulatory control, as well as the neuronal structures that orchestrate its allocation (Munakata et al., 2011; Holroyd & Yeung, 2012; Cavanagh & Frank, 2014; Shenhav et al., 2017). In the present chapter, we will demonstrate how these developments can inform and advance neuroscientific research on self-regulatory control in sports. In part one of this chapter, we will follow recent mechanistic approaches and conceptualize self-regulatory control as a reward-based decision. Specifically, we introduce the expected value of control (EVC) theory (Shenhav et al., 2013) as a mathematically explicit framework that provides a value-based computational expression for the allocation of self-regulatory control and that specifies the mechanistic foundation of self-regulatory control.

In part two of the chapter, we will discuss recent technological advancements that have enabled neuroscientific research even during full body movements, an important prerequisite for investigating neural processes during sports performance (Ekkekakis, 2009a; Perrey & Besson, 2018). This has enabled researchers to satisfy recent calls to investigate the "sporting brain" (Walsh, 2014, R859) and to examine whether findings from basic cognitive neuroscience can be applied to the field of sports and exercise. We will summarize neuroscientific research in sports through the lens of self-regulatory control, with a specific focus on functional near-infrared spectroscopy (fNIRS) as a neuroscientific method that appears to be particularly suited for research in sports.

Computational and Mechanistic Operating Principles of Self-Regulatory Control

In this section, we will elaborate on the concept of mental effort and why it is central to the definition of self-regulatory control. We will then introduce a mechanistic account of the processes that guide the allocation of mental effort in the service of self-regulatory control, based on the EVC theory.

Self-Regulatory Control Is Effortful and Costly

In keeping with this chapter's focus on physical performance, we will use an analogy from physical effort to aid the definition of mental effort. Assume, for instance, that a marathon runner is able to run a marathon in 02:08:00h (i.e. capacity). To qualify for the Olympics, he needs to run it in < 02:11:30h (i.e. task characteristics). Physical effort is what mediates between his running capacity and the time that is required for qualification, on the one hand, and the ultimately achieved finishing time, on the other hand. Likewise, mental effort can be conceptualized as that which "mediates between (a) the characteristics of a target task and the subject's available information-processing capacity and (b) the fidelity of the information-processing operations actually performed, as reflected in task performance" (Shenhav et al., 2017, p. 100). Self-regulatory control is believed to be the force through which mental effort is exerted (Shenhav et al., 2017). Tasks vary in the degree to which they can rely on controlled vs. more automatic processes (Posner & Snyder 1975; Schneider & Shiffrin, 1977). One of the major factors that can determine the control requirements of a task is the extent to which pursuing the task goal requires the individal to overcome more automatic (e.g. default) responses (Cohen, Dunbar, & Mcclelland, 1990; Shenhav et al., 2013). To illustrate, when the runner is hurting, the automatic response tendency would be to stop. However, to qualify for the Olympics, this response needs to be controlled. Thus, although in this case the behavioral output relies on physical effort (e.g. engagement of locomotor muscles), self-regulatory control (and therefore mental effort) is needed to override the automatic response tendency of stopping or slowing down.

Although it is often crucial for success, people tend to avoid engaging in self-regulatory control because the mental effort required (and associated experiences of fatigue and frustration; Wolff, Sieber, Bieleke, & Englert, 2019) is experienced as aversive (Kool & Botvinick, 2018; Westbrook & Braver, 2015; see also Box 1). Thus, applying self-regulatory control appears to carry an intrinsic disutility (Kool & Botvinick, 2018). Accordingly, people tend to mobilize effort only if the goal is subjectively worth it (Gendolla & Richter, 2010) and not to a greater degree than is warranted by the difficulty of a given task (Wright, Mlynski, & Carbajal, 2019). This indicates an aversion against mobilizing more effort than necessary (Richter, Gendolla, & Wright, 2016). Returning to the example of the marathon runner: If the *only* goal is to qualify for the Olympics (i.e. no other incentives like winning the race or beating a rival are present), he should only run as fast as needed to qualify. This reasoning is in accordance with a large body of research showing that people try to conserve their resources when it comes to the mobilization of effort (Richter et al., 2016). Taken together, research suggests that self-regulatory control is treated as if it is costly and limited.

Box 1: When Effort Adds Value

Mental effort plays a key role in our understanding of self-regulatory control. An overwhelming body of research shows that mental effort produces costs that people generally try to avoid (Kurzban, 2016). Sports seems to be at odds with this *law of least effort* (Hull, 1943). For most sports, especially endurance sports like running or cycling, physical and mental efforts are not only instrumental to reaching one's goal; they often constitute the goal itself – that is, effort is often *central* to the sporting experience. To use the words of multiple Tour de France champion Greg LeMond: "it never gets easier, you just go faster" (Missel, 2017, para. 7). Yet, every year, millions of people pay considerable amounts of money to participate in running competitions with neither the chance nor the expectation of winning (Andersen, 2020). For most of them, the cost of participating in such events – ranging from entry fees and costs for equipment and travel to the opportunity costs of training and the risk of suffering an injury (Maxcy, Wicker, & Prinz, 2019) – seems disproportionate compared to what they receive in return. This makes it difficult to explain why so many people freely choose to participate in sporting competitions. Interestingly, one experience recreational athletes appear to seek out in a marathon, or an ultra-marathon might simply be the experience of pushing through perceived boundaries and still keep going (Finn, 2018).

As this example illustrates, some people engage in sports *because* it requires mental and/or physical effort and not despite of it (Loewenstein, 1999). This implies that effort (mental or physical) may in a certain case add value to an activity. From the perspective of the EVC framework, effort can add value in two different ways (Inzlicht, Shenhav, & Olivola, 2018): First, it might amplify the value of the expected outcome. Thus, success is more rewarding when it is achieved with high effort and failure feels worse when one had invested a lot of effort. In the words of Brazilian soccer legend Pelé: "The more difficult the victory, the greater the happiness in winning" (Keville, 2015, para. 3). Second, effort itself might be intrinsically rewarding irrespective of the outcome. One possible explanation for this might be a phenomenon referred to as learned industriousness (Eisenberger, 1992): If athletes learn to directly associate effort with reward, the effort assumes the role of a secondary reinforcer. Indeed, athletes and their coaches often emphasize the need to make an effort, irrespective of the direct outcome. For example, golfer Arnold Palmer stressed that one should "always make a total effort, even when the odds are against you" (Park, 2013, para. 38). And the value of effort is very clearly expressed by Eddy Merckx – considered by many to be the greatest cyclist of all time – who stated that "when it's hurting you, that's when you can make a difference" (Alé La Merckx, 2018, para. 5). The paradoxical nature of effort as both costly and rewarding is not limited to experiences in the sports setting, and has recently started to receive research interest (for a review, please see Inzlicht et al., 2018), although comparatively little research has yet been conducted to investigate the ways in which effort adds value. Sports seems to be a domain where effort is held in very high regard and treated as an incentive in its own right. The field of sports and exercise thus serves as a prime target for the investigation of this effort paradox.

Various explanations have been proposed for *why* the capacity for self-regulatory control might be limited. For example, resource-based accounts suggest that self-regulatory control depletes a metabolic resource (see Chapter 8, this volume; Baumeister, Bratslavsky, Muraven, & Tice, 1998). In contrast to such structural limitations in the capacity for self-regulatory control, other accounts focus on the computational properties of the neuronal systems that govern control. From this perspective, exertion of self-regulatory control is perceived as effortful not because a resource has been depleted but to index the costs that arise in tasks that compete for the same neural representations (Feng et al., 2014; Musslick et al., 2016). This explanation can be transferred to the context of sports and exercise: If an athlete is applying self-regulatory control to perform at her best in a cycling time trial, performance in concurrent tasks that rely on control will likely be impaired (if performance is not impaired, then the task should at least be perceived as more effortful). For example, while the athlete has to control the impulse to go slower, a concurrent task that also relies on impulse control (e.g. to control one's emotional expression) would require self-regulatory control to manage both tasks concurrently, thereby leading to impaired performance and/or increased sensation of mental effort. Tentative support for this interpretation in the context of physical performance comes from the occurrence of cognitive motor interference that already occurs in simple motor tasks like walking (for a meta-analysis, please see Al-Yahya et al., 2011). Likewise, impaired cognitive performance in dual-task situations has been found in diverse sport settings such as table tennis (Schaefer & Scornaienchi, 2020), swimming (Stets, Smith, & Helton, 2019), or climbing (Darling & Helton, 2014). For example, when participants were asked to perform an n-back task while returning table tennis balls from a ball machine, this led to impaired n-back performance (Schaefer & Scornaienchi, 2020). Interestingly, this dual-task cost was signifcantly more pronounced in novices than in experts, which might point toward more automated processing of demands in experts (Schneider & Shiffrin, 1977).

While the exact reasons that limit the brain's capacity for self-regulatory control are still debated (Kurzban et al., 2013; Shenhav et al., 2017), it appears to be clear that people invest mental effort sparingly and treat its mobilization as if the capacity for control is limited. Consequently, researchers have tried to understand the factors that determine how people choose to allocate self-regulatory control.

The Expected Value of Control Theory

Research on self-regulatory control in sports has long been dominated by resource-based conceptualizations of self-regulatory control (for an overview, please see Englert, 2016). Recently, these resource-based accounts have been challenged empirically (Carter & McCullough, 2014; Hagger, Chatzisarantis et al., 2016; Wolff, Baumann, & Englert, 2018) as well as conceptually (Inzlicht et al., 2014; Kurzban et al., 2013), and recent theoretical accounts now converge toward conceptualizing self-regulatory control as some form of reward-based choice (Berkman, Hutcherson, Livingston, Kahn, & Inzlicht, 2017; Inzlicht et al., 2014; Kool & Botvinick, 2014; Kurzban et al., 2013; Wolff & Martarelli, 2020). However, these theoretical accounts have so far rarely been adopted in research on self-regulatory control in sports and exercise. This is puzzling because reward-based accounts play an important role in research on motor control (Körding, Fukunaga, Howard, Ingram, & Wolpert, 2004; Manohar et al., 2015; Morel, Ulbrich, & Gail, 2017; Shadmehr, Huang, & Ahmed, 2016), which is fundamental to any sporting activity. Here, we present the EVC theory as one such reward-based model of self-regulatory control (Shenhav et al., 2013, Shenhav, Cohen, & Botvinick, 2016). The EVC theory synthesizes core concepts of various control theories into an integrative framework that specifies the computational properties as well as the mechanistic underpinnings of self-regulatory control (for an illustration, please see Figure 11.1a).

Computation of the EVC

According to the EVC theory, people allocate self-regulatory control in a way that maximizes the expected payoff while minimizing the mental effort that is required. These collectively determine the expected value of a given control allocation (EVC). The outcome of this cost-benefit analysis (the EVCs of candidate control configurations) determines which task(s) are most worth allocating self-regulatory control to, and how much control should be allocated (Figure 1b). Equation 11.1 formalizes how the EVC is calculated (Shenhav et al., 2013):

$$\mathrm{EVC}(\mathrm{signal},\ \mathrm{state}) = \left[\sum_i \Pr(\mathrm{outcome}_i \mid \mathrm{signal},\ \mathrm{state}) \cdot \mathrm{Value}(\mathrm{outcome}_i)\right] - \mathrm{Cost}(\mathrm{signal}) \quad (11.1)$$

The situation (denoted as *state*) a person applies control to is shaped by internal and external states the person finds itself in. For example, a marathon runner might face a strong headwind (external) and also have some lingering doubts regarding her stamina over the full distance (internal). In any given state, different potential control signals of varying vigor can be specified (Figure 1a). The marathon runner's overall goal might be to qualify for the Olympic Games. To do so, she has to complete the distance faster than a set target time and this is likely to be very hard. Thus, while running, she has to control the automatic impulse to ease up and instead stick with her target running pace. Over the course of the competition, the effort required for sticking to the race pace will rise due to fatigue that is setting in.[2] The athlete might reason that she can only consistently resist the impulse to ease up as long as perceived physical exertion does not go beyond 80% of her maximum in the first half of the race, to avoid what runners refer to as "hitting the wall". Further, this might allow her to increase the intensity late in the race and produce a final spurt if needed. Control signals therefore vary in *identity* (e.g. "run" and "inhibit inclination to ease up") and in *intensity* (e.g. for running this might vary from 0% to 100% of maximum effort).

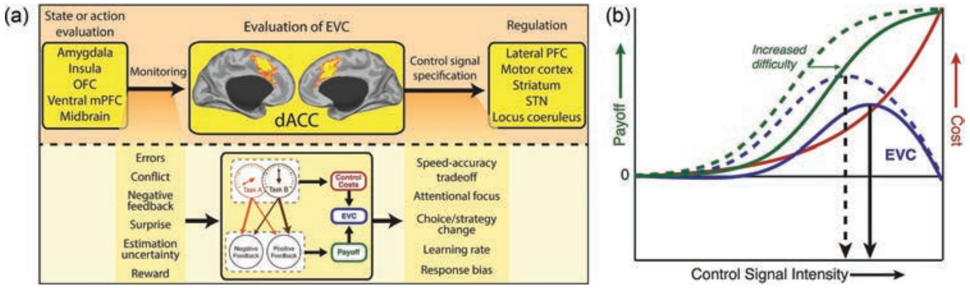

FIGURE 11.1 The Expected Value of Control (EVC) theory. (a) The EVC theory predicts that dACC integrates information about relevant incentives and task demand to determine the overall EVC of potential control configurations. On the basis of this EVC calculation, the dACC specified the optimal (EVC-maximizing) set of control signals, and signals these to relevant downstream regions which guide the execution of these control signals. Adapted from Shenhav et al. (2016). (b) Within the EVC framework, increases in the expected difficulty of a task (bold vs. dashed lines) will often lead to the prediction that additional control is needed to achieve a given level of payoff (green curves). This expected payoff is discounted by the expected cost (red curves) of exerting the necessary mental effort (control intensity), determining the overall EVC of that control intensity (blue curves). The maximum of this EVC curve determines the optimal control setting (black arrows). Adapted from Shenhav et al. (2013).

EVC can be calculated for any available configuration of control signal identities and intensities. As can be seen in the right part of Equation 11.1, the EVC of any given control signal configuration comprises two main components: An expected payoff and a cost. The expected payoff is determined by considering potential future outcomes that are relevant to one's task – including potential praise and pride that come with success, and potential admonishment and embarrasment that come with failure – and weighing these potential outcomes by the likelihood that they would occur under a particular control setting. For instance, as someone increases their physical exertion, they increase the likelihood of success (and attendant positive incentives) and decrease the likelihood of failure (and attendant negative incentives). However, increasing physical exertion is also costly: All things being equal, exercise at very high intensities (e.g. 95% of maximum) generally elicits a more negative affective response than exercise at lower intensity (e.g. 80% of maximum) (Ekkekakis, Parfitt, & Petruzzello, 2011; Roloff et al., 2020). This cost is accounted for by the final term in the equation, which trades off against the expected payoffs, discounting the potential positive outcomes that accrue from increasing intensities of physical exertion. As can be inferred from Equation 11.2,

$$\text{Value}(\text{outcome}) = \text{ImmediateReward}(\text{outcome}) + \gamma \max_i \left\lfloor \text{EVC}\left(\text{signal}_i, \text{outcome}\right) \right\rfloor \tag{11.2}$$

the outcome value reflects the value of the immediate change that occurs due to applying self-regulatory control (denoted as *ImmediateReward*), as well as accounting for the EVC of future control signals that are feasible based on the chosen signal (expressed as a maximization of EVC taking into account feasible control signals). This recursive definition of outcome value is important, because it highlights that the chosen EVC is not only maximized based on the immediate reward but also on the value of more distal states. For example, the *ImmediateReward* of running only at 80% of maximum exertion might in fact be negative, because the athlete will run slower than at 95% of maximum exertion and this might lead to the athlete being dropped from the lead group. However, specifying a control signal of such (comparably) little intensity might allow this athlete to maintain the target pace over the full distance, better positioning her to achieve her goal of qualifying for the Olympics. Thus, the value of an outcome is the sum of proximal and distal changes that are brought about by the specified control signal. However, as humans prefer immediate rewards over future rewards (Critchfield & Kollins, 2001; Frederick, Loewenstein, & O'Donoghue 2002), a discounting factor γ is introduced to discount the impact of future EVCs in determining the overall value of the outcome that is achieved by the current control signal. This discounting factor is a variable ranging from zero to one to account for intra- and inter-individual differences in discounting. Interestingly, a large-scale analysis of 1.7 million recreational runners' pacing strategies showed that relatively fast starts and fast finishes were both predictors of poor overall performance (Smyth, 2018). From an EVC perspective, this might be attributed to a suboptimal discounting factor (i.e. by over- or underweighting future EVCs).

EVC theory proposes that the human control system performs these calculations, compares the EVCs for different control signals, and then selects the one where EVC is maximized. Recent work has validated these predictions by showing that simulated agents that are designed to maximize their EVC adjust their task performance based on the incentives and task demands in similar ways as has been observed in empirical studies of motivation and self-regulatory control (Musslick, Shenhav, Botvinick, & Cohen, 2015; Lieder, Shenhav, Musslick, & Griffiths, 2018; Bustamante, Lieder, Musslick, Shenhav, & Cohen, 2020).

Mechanistic Underpinnings of Self-Regulatory Control

In addition to specifying the computational properties that underlie self-regulatory control, EVC theory also specifies candidate neural structures that mediate different processes involved in allocating and adjusting self-regulatory control (Shenhav et al., 2013). EVC theory delineates three core functions: *Specification*, *Regulation*, and *Monitoring.* Self-regulated control consists of a continuous loop in which control signals are specified, applied, and the outcome is monitored to assess whether the control signal needs to be adjusted. EVC theory proposes that the dorsal anterior cingulate cortex (dACC) sits at the interface of monitoring and specification, whereas regulation is primarily mediated by regions downstream of dACC that execute specific types of self-regulatory control, such as the lateral prefrontal cortex (lPFC). The functions EVC theory ascribes to dACC and lPFC are well supported by a large body of research (for reviews, please see Badre, 2008; Heatherton, 2011; Miller & Cohen, 2001; Shenhav et al., 2013).

Adaptive control thus requires an individual to monitor their current state and the value of potential future outcomes that might be attained with available control. Such monitoring is intuitively important in sports and exercise: For athletic success, it is crucial to perform at (or close to) one's physiological limit without overextending oneself. Thus, athletes need to closely monitor their current state (e.g. pace) and flexibly adjust their behavior if this state is at odds (e.g. too slow) with the goal they aim to achieve (e.g. qualifying for the Olympics). As we have outlined above, these adjustments are made in order to invest self-regulatory control efficiently. In regard to control signal intensity, this means that the EVC is maximized when the intensity is as low as possible while still enabling goal attainment. For a marathon runner trying to qualify for the Olympic games, this implies that she will only invest the amount of mental effort needed to *just* achieve the qualifying time.[3] However, the emphasis on not wasting resources might sometimes lead to the specification of a control signal intensity that is too low. The marathon runner might, for example, follow a conservative pacing strategy (e.g. negative split times) but if she is unable to make up for lost time later in the race, she might eventually miss the qualification time. Thus, selecting an optimal pacing strategy[4] is crucial to athletic success (Abbiss & Laursen, 2008).

Research shows that the dACC monitors various indicators of state information (e.g. information that is relevant for control signal identity and intensity specification) in a fashion that is likely to subserve the computation of the EVC (Shenhav et al., 2013, 2016). The plethora of different bits of information that might be monitored during exercise is provided by various brain areas (Figure 11.1a): Research indicates that interoceptive information (e.g. heavy breathing, locomotor fatigue) is processed by the insula (e.g. Gehrlach et al., 2019; Livneh et al., 2020), rewards (e.g. joy of running fast, achieving the qualification) are encoded by the ventromedial PFC (Gläscher, Hampton, & O'Doherty, 2009; Strait, Blanchard, & Hayden, 2014), and the amygdala may signal negative affective states such as fear of failure and exercise-induced pain (Neugebauer, Li, Bird, & Han, 2004). The dACC monitors signals from these and other structures as information for computing the EVC. The monitored information then needs to be integrated to specify a control signal that maximizes EVC. Research indicates that the specification and, if required, adjustment of the control signal is performed by the dACC (Cavanagh & Frank, 2014; Phillips, Johnston, & Everling, 2011; Shackman et al., 2011; Ullsperger, Danielmeier, & Jocham, 2014).

The EVC-maximizing control signal is then relayed to relevant downstream regions for execution. For many forms of self-regulatory control (e.g. controlling the impulse to get into a shoving match with an opponent after a foul; pulling oneself together to go out to train in the pouring rain), the region that appears to be most critical for this execution is lPFC (Miller & Cohen, 2001). The dACC and lPFC share strong reciprocal connections (Heilbronner & Hayden, 2016), and tend to co-activate in most control-demanding tasks (Duncan, 2010; Niendam et al., 2012). The strength of this dACC-lPFC co-activation has been shown to be predictive of good self-regulatory control, for instance improved performance on a rotation-letter task (Kondo, Osaka, & Osaka, 2004). Of particular relevance to our central proposition in this chapter, namely that self-regulatory control is important for controlling sports performance, lPFC is also well-connected to structures in the motor system, including the premotor cortex, which is, in turn, extensively connected to areas of relevance for direct motor output, like the primary motor cortex and the basal ganglia (Miller & Cohen, 2001). The connections between PFC and the basal ganglia are of interest in the domain of sports because the basal ganglia are believed to play a key role in generating central fatigue. Physical performance depends strongly on the capacity to tolerate fatigue, which may depend on top-down regulation of striatum by prefrontal cortical structures (Chaudhuri & Behan, 2000).

Cortical Activity during Self-Regulated Sports Performance

The bulk of the neuroscientific evidence that we have presented above stems from animal studies or human imaging studies that were performed under conditions that do not mirror real-world situations. Although it has been proposed that sports might be the "brain's biggest challenge" (Walsh, 2014, p. 859), research on the brain *while* a person is doing sport is still comparably scarce (Perrey & Besson, 2018). However, this is rapidly changing. Over the last 25 years, researchers have become increasingly capable of and interested in the investigation of the sporting brain (Figure 11.2a). Of particular interest for the present chapter is the drastic uptick in research on the brain *during* sports. To illustrate this, we conducted a web of science search that combined the topics "sports" and "brain" and that also required the word "during" in the title of published articles (Figure 11.2b). This search returned only eight hits during the 20-year time span from 1995 and 2004. In the following 15 years, another 171 papers have been published and in 2019 alone, 26 papers fulfilled these search criteria. Thus, while this literature search is by no means comprehensive, it seems safe to say that research on the sporting brain is garnering momentum.

One likely reason for this momentum is the introduction of more portable and more robust technologies for measuring cortical processes during sports. Particularly, fNIRS and, to a lesser extent, electroencephalography (EEG) have been increasingly used in the sports context in recent years (Perrey & Besson, 2018). While most readers are likely familiar with the basic operating principles of EEG, fNIRS is a comparatively recent neuroimaging technology whose specific advantages and disadvantages are less well known. Thus, prior to summarizing the literature on self-regulated sports performance, we will briefly describe the operating principles of fNIRS.

Measuring the Sporting Brain: fNIRS

Neuronal activity consumes energy and local changes in neuronal activity cause local increases in oxygenated blood. In the blood, oxygen is transported by the chromophore hemoglobin, which occurs in an oxygenated (HbO) and a deoxygenated (HbR) variant. Like functional magnetic resonance imaging, fNIRS tracks the resultant change of balance between HbO and HbR as a marker for neuronal activity. fNIRS uses NIR light to measure

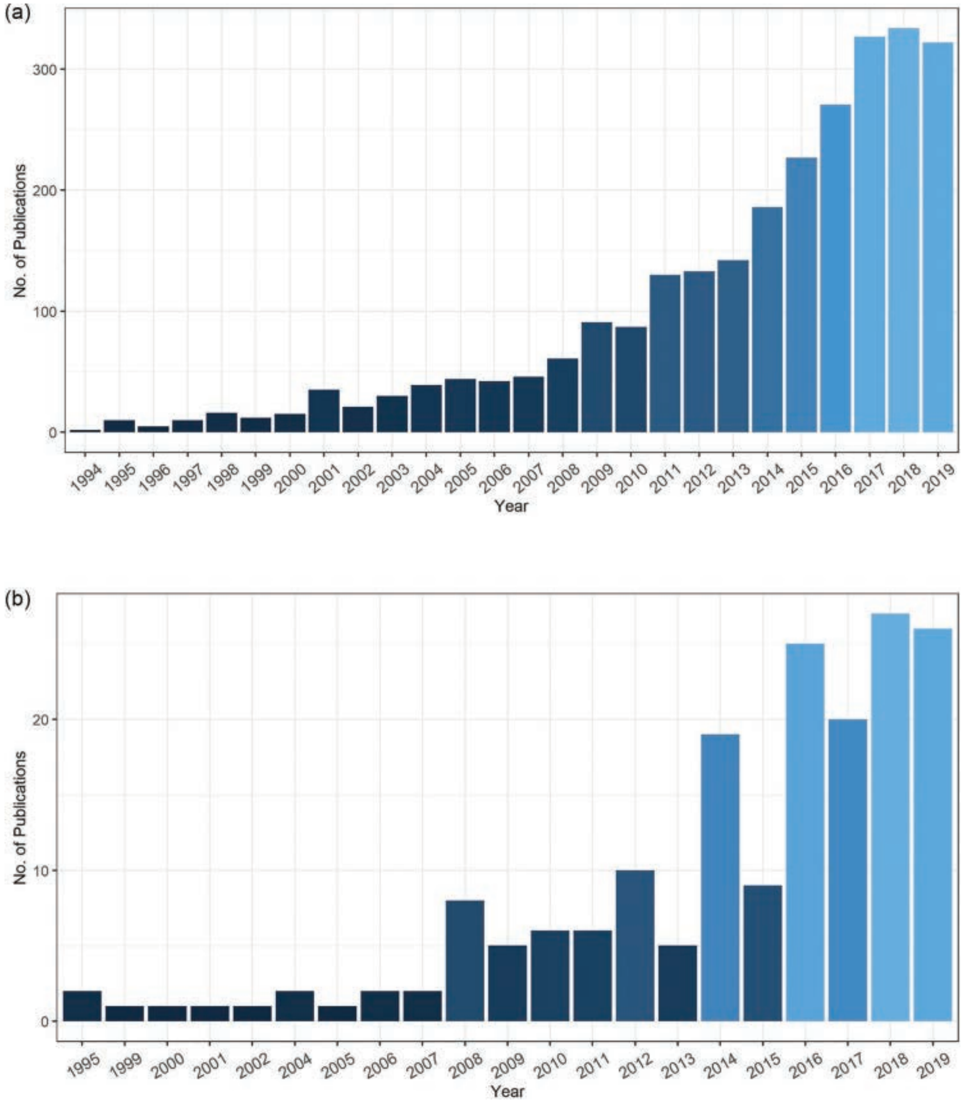

FIGURE 11.2 Web of Science Search (per January 2020) on the publication trends for neuroscientific research in regard to sports. Panel (a) illustrates the broad increase in research into the sporting brain (search terms: Sports AND brain). Panel (b) only includes papers that additionally have the term "during" in their title, to illustrate the relative scarcity and the increase in research on brain activity while people do sports.

changes in HbO and HbR non-invasively in the cerebral cortex (Scholkmann et al., 2014). To do so, NIR light emitting senders and NIR light sensitive detectors are placed on the participants' scalp (see Figure 11.3a), according to a predefined setup (see Figure 11.3b) that is designed to capture regions of interest (see Figure 11.3c). The NIR light spectrum ranges from 700nm to 1400nm and within this spectrum, Hb0 absorbs substantially more light at ranges above 830nm and HbR absorbs more light at 760nm (Ekkekakis, 2009). Capitalizing on this differentiation in absorption spectra, fNIRS senders emit light in two wavelengths

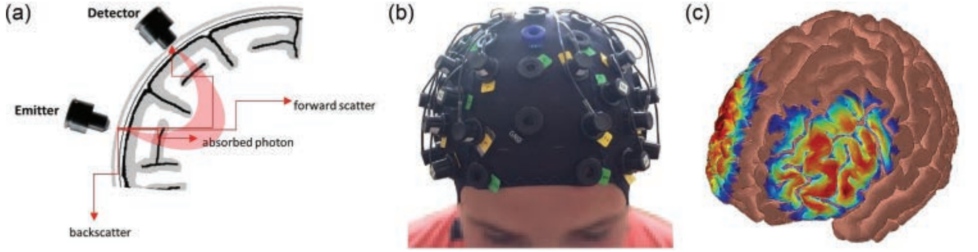

FIGURE 11.3 fNIRS Setup: (a) Schematic illustration of NIR light pathway between emitter and detector, (b) illustration of fNIRS montage designed to capture oxygenation changes in the dorsolateral prefrontal cortex, and (c) illustration of the sensitivity profile for capturing specific regions based on the montage designed to capture changes in the dorsolateral prefrontal cortex.

(e.g. 760nm and 850nm) and detectors nearby[5] track how much of the emitted NIR light leaves the scalp. Relative changes in the detected light are then analyzed as indicators for Hb0/HbR fluctuations, which serve as a proxy for change in cortical activity (Ferrari, Mottola, & Quaresima, 2004). For assessing cortical activity during sports, fNIRS has been regarded as the best suited technology that is currently available (Ekkekakis, 2009; Herold, Wiegel, Scholkmann, & Müller, 2018). Most important for sports settings, fNIRS is comparatively robust to motion artifacts, has acceptable temporal resolution and – due to its portability – can be used in settings with high external validity (Ekkekakis, 2009; Strangman, Culver, Thompson, & Boas, 2002). fNIRS has even been used to measure activity in the motor cortex in freely moving participants (Piper et al., 2014). Specifically, these authors were interested in whether hemodynamic changes that originate from hand movement could be captured while participants were sitting on a bike in decreasingly less-controlled conditions: Participants either sat on a stationary bike without pedalling, were pedalling on a stationary bike, or were freely riding their bike outside. Although an increase in external validity was acompanied by an increase in motion artifacts, the hemodynamic response that was triggered by the hand movement in the contralateral motor cortex could still be recovered (Piper et al., 2014). In addition, compared to technologies like fMRI, fNIRS is cheap to acquire and produces little running costs (Wolff, 2017). Finally, compared to fMRI and EEG, fNIRS measurements tend to be perceived as less invasive and aversive by participants (Cutini & Brigadoi, 2014). Constraining its applicability, fNIRS can only be used to reliably measure oxygenation changes in superficial cortical areas, whereas limited depth penetration prevents the measurement of deep brain areas (Ferrari & Quaresima, 2012). With respect to the key cortical structures involved in self-regulatory control, fNIRS is therefore well-suited for monitoring activity in the lPFC but not in the dACC. Thus, most of the literature we review below has focused on lPFC, whereas research on dACC is scarce. Further, although the spatial resolution of current fNIRS devices is ≤1 cm (Ferrari & Quaresima, 2012), fNIRS is less accurate than the spatial resolution that can be achieved with fMRI. Thus, for research that requires very precise localization or is concerned with subcortical areas, fNIRS is less well-suited. However, to assess group-level effects (e.g. comparing groups of trained and non-trained participants), fNIRS has proven to be a reliable neuroimaging technique (Scholkmann et al., 2014) and has accordingly gained in popularity in the domain of sports in recent years (for a more comprehensive introduction of fNIRS to research in sports, please see Ekkekakis, 2009; Perrey & Besson, 2018).

Brain Activity during Self-Regulated Sports Performance

Only recently have researchers started to explicitly conceptualize sports performance as a self-regulatory challenge (Marcora, 2008; McCormick, Meijen, Anstiss, & Jones, 2018; Pageaux, 2014; Wolff, Bieleke et al., 2018). Consequently, many early studies on the brain's role in exercising have focused on explaining exercise-induced activation changes in the brain from a physiological point of view (Rooks, Thom, McCully, & Dishman, 2010; Secher, Seifert, & Van Lieshout, 2008). For example, a systematic review of 25 fNIRS studies that were published between 1999 and 2009 concluded that with increasing exercise intensity, activation in the lPFC increased until it dropped at very high intensities (Rooks et al., 2010). One interpretation of this pattern suggested that this reflects an automatic physiological response that is initiated to prevent harmful threats to bodily homeostasis when exercise intensity is dangerously high[6] (Noakes, Peltonen, & Rusko, 2001; Rooks et al., 2010). However, rather than reflecting an obligatory response to a physiological limitation, these same findings could reflect the endpoint of a decision process, which weighs the value of applying further effort against the cost of this effort. Specifically, these changes in lPFC activity could be conceptualized in terms of variability in self-regulatory control: During an incremental exercise protocol (for example, riding on a cycling ergometer until exhaustion), the self-regulatory demands rise steadily because the task becomes increasingly more challenging and participants have to fight the impulse to quit. From an EVC perspective, to apply the self-regulatory control that is required to continue in spite of these aversive sensations, the control signal intensity needs to be constantly upregulated by the dACC and more regulative control needs to be applied by the lPFC (Figure 1b). Thus, if one understands an incremental exercise task as a self-regulatory challenge (Wolff, Bieleke et al., 2018), an increase in lPFC oxygenation is expected over time. This assumption is supported by findings of recent fNIRS studies. For example, in one study, lPFC oxygenation during a 30-minute cycling task increased as a function of duration and intensity in a sample of regularly exercising athletes (Giles et al., 2014), probably accompanying increasing levels of self-regulatory control. Similarly, when participants had to produce 50% of their maximum voluntary handgrip force, this self-control-demanding performance was associated with higher lPFC oxygenation, compared to a condition when only 30% of maximum force was required (Guo, Ma, & Chen, 2019). However, at a certain exercise intensity, the cost of exerting the level of self-regulatory control that would be required to continue with the task (e.g. controlling the impulse to stop) will exceed the expected payoff from the task (e.g. from receiving course credit for participating in an experiment). From an EVC perspective, this shift in EVC should motivate the person to reduce the intensity of self-regulatory control they are applying (Frömer, Lin, Dean Wolf, Inzlicht, & Shenhav, 2020), leading to a drop in PFC oxygenation and subsequent task termination. Indeed, an EEG study where cyclists completed an incremental cycling test showed that after an initial rise in PFC activation, a significant drop was observed after the respiratory compensation point (Robertson & Marino, 2015). A similar result could be observed in an fNIRS study with trained cyclists: In a progressive maximal cycling exercise to exhaustion, lPFC oxygenation increased in the first part of the cycling task but decreased significantly prior to task termination (Rupp & Perrey, 2008). Interestingly, a drop in lPFC activation does not lead to immediate task termination but participants continue for some more time[7] (Wolff, Bieleke et al., 2018). One possible explanation for this finding is that exercisers specify control signal intensities that have a built-in anticipatory buffer to be prepared in case exercise gets more demanding. The quick drop in lPFC oxygenation that has been observed prior to task termination might then reflect the dropping of EVC until the specified control intensity does not suffice to continue with the task. This is in line with research showing that a certain activity level in lPFC is required to engage in an effortful task (Hosking, Cocker, & Winstanley, 2016). In addition, this pattern might also reflect some form of task inertia: Changing a self-regulatory control signal produces costs in itself and these reconfiguration costs can deter people from switching to another activity (Monsell, 2003; Musslick, et al., 2018).

Trained Athletes Process Exercise-Induced Demands More Efficiently

Some studies have compared exercise-induced cortical activation changes in trained athletes and untrained participants (Ludyga, Gronwald, & Hottenrott, 2016; Rooks et al., 2010; Seidel, Carius, Roediger, Rumpf, & Ragert, 2019). Interestingly, in the systematic review by Rooks et al. (2010), the drop in PFC oxygenation prior to task termination has been observed in untrained participants but not in trained athletes. Trained athletes have extensive experience in adjusting their mental and physical efforts as a function of the challenge at hand, by interpreting sensations from their body. It is therefore conceivable that trained athletes are better able to anticipate expected control costs and the likelihood of potential future rewards, resulting in better-calibrated estimates of EVC. Behavioral data in support of this hypothesis comes from an analysis of the pacing profiles of 190,228 finishers of the New York Marathon. Here, the variability in 5km split times was substantially lower for top runners compared to less successful runners (Santos-Lozano, Collado, Foster, Lucia, & Garatachea, 2014).

Further, the *neural efficiency hypothesis* suggests that experts are more efficient in using their cortical resources for performing mental operations (Dunst et al., 2014). For example, more intelligent individuals perform cognitive tasks with lower brain activation, indicating the investment of less mental effort (Dunst et al., 2014). Applying this finding to the sports setting, trained athletes exhibit a less pronounced increase in PFC oxygenation during incremental exercise than untrained participants (Rooks et al., 2010). This is also in line with research showing that participants who score highly on a self-report measure of trait self-regulatory control display a less steep increase in lPFC oxygenation over the course of a fatiguing static muscular endurance task (Wolff, Schüler et al., 2019). This might point toward a more efficient use of processing resources. Interestingly, when task demands were standardized with respect to participants' individual capability (e.g. 60% of one's maximum power output on a cycling task), these expert-novice differences in brain oxygenation have not been observed (Dunst et al., 2014; Seidel et al., 2019). Adding to evidence from fNIRS studies, a recent study used EEG to measure the α-wave/β-wave ratio as a marker for neural efficiency,[8] and compared well-trained to less well-trained cyclists (Ludyga et al., 2016). Higher α/β ratios were observed over frontal sites in well-trained cyclists, supporting the neural efficiency hypotheses. Although they do not explicitly refer to self-regulatory control, the authors suggest that superior athletes might be better at inhibiting task-irrelevant cognitive processes. Going back to the proposition that the control demands of a task vary as a function of its automaticity and of its reliance on shared local processing pathways, these findings fit well into the theorizing on self-regulatory control we present here. Compared to non-athletes, trained athletes have been repeatedly exposed to the multitude of self-regulatory challenges one faces during exercise (McCormick, Meijen, & Marcora, 2016) and this experience likely leads to a greater automatization in the processing of these regulatory demands.[9] Thus, it is conceivable that trained athletes process self-regulatory demands during sports more efficiently and specify more realistic EVCs that are more accurately matched to the current task demands.

Psychological Manipulations Affect Cortical Changes during Sports

Recently, researchers have started to investigate the effect psychological manipulations have on effort-related sensations and on cortical activity while participants perform sports. For example, one way to reduce the self-regulatory costs of an action is to make that action more automatic (Bayer, Achtziger, Gollwitzer, & Moskowitz, 2009; Gollwitzer, 1999; see also Chapter 12 for an in-depth discussion of the topic). Accordingly, researchers have investigated the effects of self-regulatory strategies that aim to automatize behavior (Bieleke & Wolff, 2017; Wolff, Bieleke et al., 2018). Interestingly, while such a self-regulatory intervention did not lead to improved performance in a static muscular endurance task, participants were able to perform the task with a less pronounced increase in lPFC activation, indicating that the task had been performed in a neuronally more efficient way (Wolff, Bieleke et al., 2018). Another study used the multi-action plan (MAP) model (Robazza, Bertollo, Filho, Hanin, & Bortoli, 2016) to derive functional and dysfunctional pacing strategies and assess their effect on cycling performance and on EEG coherence (as a measure of functional connectivity; Di Fronso et al., 2018). When applying the dysfunctional strategy, participants were asked to focus their attention internally on feelings of muscle exertion. According to the MAP model, this internal focus on aversive sensations is detrimental to performance because it enhances perceived fatigue. Indeed, applying a dysfunctional strategy was associated with higher EEG coherence at high intensities, potentially indicating excessive attentional focus on muscle exertion (Di Fronso et al., 2018). Interpreting this finding from the perspective of self-regulatory control, adopting a dysfunctional pacing strategy might have resulted in inefficient or excessive allocation of self-regulatory control. In line with this argument, low coherence was observed when participants adopted a functional pacing strategy. Other researchers induced heat stress to alter the costs of exercise. Here, changes in frontal areas of the brain were observed, again likely reflecting changes in the "capacity to sustain mental readiness and arousal" (Périard, Pauw, Zanow, & Racinais, 2018, p. 1).

We are aware of only one study that has investigated the influence of experimentally manipulating the value of self-regulatory control on lPFC oxygenation and physical performance: In this fNIRS study, participants performed an isometric knee extension task until voluntary task termination on two occasions (Giboin, Gruber, Schüler, & Wolff, 2019). In the control group, participants received the same reward per minute in both sessions. Participants in the experimental group received 50% less reward per minute in Session 2 compared to Session 1, but were also compensated with an upfront payment that was calibrated so they would earn exactly the same total reward in the second session if they performed the task for the same duration as in the first session. Previous research has shown that this form of *income-compensated wage decrease* (ICWD) reduces the willingness to invest self-regulatory control during that second session (Kool & Botvinick, 2014). Consistent with this finding, lPFC oxygenation was lower in the ICWD group in the second session. However, counter to expectations, this decrease was not associated with poorer performance, as would be expected if this group was exerting less mental effort. Instead, the ICWD group in fact persisted *longer* in the knee extension task. One potential explanation for this is that the unexpected performance improvement we observed was achieved by a more efficient execution of the task. Indeed, post-hoc analyses revealed that participants in the ICWD group deviated less from the target force that had to be produced, thereby executing the task with less force than the control group. There were no group differences in perceived effort and in a host of objective markers of physiological exertion (maximal voluntary contraction, voluntary activation, potentiated twitch at rest, and electromyographic activity). Collectively, these findings suggest that, under certain conditions, increasing the value of self-regulatory control might lead to an overexertion of physiological resources and thereby produce more costs and hamper performance. Interestingly, this aligns well with recent work showing that predictions people make from previous reward learning can cause them to overexert self-regulatory control (Bustamante et al., 2020).

Another branch of research has investigated the effects of providing inaccurate information regarding the duration of an endurance task on lPFC oxygenation and performance. For example, in one study, participants either knew or did not know how long a cycling task was going to last (Wingfield, Marino, & Skein, 2018). Withholding information on task duration can be understood as a manipulation that complicates the initial computation of the EVC and, subsequently, the adequate updating of EVC based on ones' current state. In line with this reasoning, participants who knew the distance invested more effort at the end of the task which was accompanied by an increase in lPFC oxygenation. In contrast, participants without this information adopted a more conservative pacing strategy and performed worse (Wingfield et al., 2018). A similar pattern was observed in a cycling study where one group of cyclists was told to cycle for ten minutes, while a second group was told to cycle for sixty minutes (Radel, Brisswalter, & Perrey, 2017). However, in the second group, the task was terminated after ten minutes as well. Participants who expected to expend their effort over a longer period of time displayed lower lPFC activation than participants who expected to cycle for only ten minutes.

Beyond Endurance Performance

The bulk of the neuroscientific literature we reviewed here concerned endurance performance. This has at least two reasons: First, the self-regulatory challenges of endurance performance are ideal to study self-regulatory control, given that the regulation of effort is at the heart of endurance sports and that perceived exertion is a fundamental sensation that characterizes the application of self-regulatory control (Wolff & Martarelli, 2020). Second, endurance performance – especially time trial cycling or static muscular endurance – allows for a comparatively controlled and stable measurement of cortical processes because the upper body is relatively stable and athletes usually do not make abrupt movements.

However, researchers have also investigated other types of sports, and their findings are compatible with those obtained in the domain of endurance performance. For example, a recent study found that a more complex juggling task (e.g. 5-ball cascade vs. 2-ball cascade) was accompanied by more oxygenation over motor areas, whereas a higher level of juggling expertise tended to be associated with less oxygenation (Carius et al., 2016). However, as activity was only monitored over motor areas it remains unclear whether the juggling task also elicited an increase in lPFC oxygenation. Considering the projections from the lPFC to premotor areas (Miller & Cohen, 2001) and the self-regulatory demands imposed by the task (i.e. attentional control for tracking ball trajectories), this would be likely. Another recent study investigated changes in oxygenation in the lPFC during a series of sprint start sequences (Wolff, Thürmer, Stadler, & Schüler, 2019). Participants were instructed to either produce a fast sprint start on the Go-signal or *not* to start on the Go-signal (no-start condition). Thus, in the condition where participants actually had to start, they had to balance the impulse to start too fast with the requirement of starting as fast as possible on the Go-signal. In the no-start condition, by contrast, participants only had to restrain the impulse of accidentally starting upon the Go-signal. From an EVC perspective, both tasks required a control signal to be specified to prevent a false start (i.e. response inhibition). However, when actual starts were required, an additional control signal needed to be specified to initiate a fast start as soon as the Go-signal occurred (i.e. action initiation). In line with this reasoning, an increase in PFC oxygenation in the period between the Set-signal and the Go-signal was observed in both conditions. However, this increase was even more pronounced in the condition where participants actually had to produce a start.

Conclusion

Taken together, the literature reviewed in this chapter demonstrates that research on the neuroscience of self-regulatory control in sports is still in its infancy but has started to gather considerable momentum (see Figure 11.2). Performance in sports is increasingly understood as a task with inherent self-regulatory demands that athletes must cope with effectively. This perspective opens the field for innovative insights from research addressing the "whys" and "whens" of successful self-regulatory control. In particular, theories that conceive of self-regulatory control as the output of a reward-based decision offer a comprehensive computational framework that allows specific predictions regarding behavior and its neurophysiological mechanisms. As we have shown, components of one such framework, the EVC theory, can be plausibly related to various aspects of performance in sports, ranging from the initial willingness to invest effort to the decision to terminate a task. Fortunately, state-of-the-art brain imaging technology nowadays permits several intriguing ways of capturing and manipulating neurophysiological processes during exercise, providing glimpses into what has been referred to as the "sporting brain" (Walsh, 2014, R859). That said, we feel it is important to call for more neuroscientific research on self-regulated sports performance because key questions have not yet been addressed comprehensively. It will be important to test the dACC's proposed role during sports performance, to analyze the effect of directly manipulating the value of self-regulatory control, to investigate the potential neurocomputational overlap between physical and mental efforts (Ritz, Frömer, & Shenhav, 2020), and to get a better understanding of the differences between elite and recreational athletes in the neuronal processing of the self-regulatory demands in sports. We believe these are exciting questions and with this chapter we hope to encourage further research on self-regulatory control in sports.

Notes

1 Sometimes also referred to as *self-control* or *cognitive control*.
2 Indeed, research indicates that perceived exertion, rather than physiological resource depletion (e.g. locomotor muscle fatigue), might be the signal that ultimately leads to the termination of an

exercise (e.g. Chapter 10, this volume; Marcora, Staiano, & Manning, 2009; Staiano, Bosio, de Morree, Rampinini, & Marcora, 2018).

3 Clearly, this example only applies to a scenario where the runners' only goal is to qualify and not to achieve a personal best or win the race.

4 Interestingly, research on pacing implies a process with a striking similarity to EVC theory, when it comes to how athletes pace themselves during competition: "intensity is regulated within the brain based on a complex algorithm involving peripheral sensory feedback and the anticipated workload remaining" (Abbiss & Laursen, 2008, p. 239).

5 Research suggests that sender-detector separations of ca. 3 cm optimize signal-to-noise ratio (Boas, Elwell, Ferrari, & Taga, 2014; Ekkekakis, 2009).

6 Indeed, research shows that task termination occurs well before physiological resources are fully depleted or bodily homeostasis is threatened (Marcora & Staiano, 2010; Morales-Alamo et al., 2015).

7 Systematic research on how long participants continue after this drop is currently lacking and represents an exciting future research question.

8 α-waves are associated with mental inactivity and β-waves have been linked with cognitive information processing. Thus, if a demanding task is performed with a high α-wave/β-wave ratio, this can be interpreted as an indicator for energy-efficient cortical functioning (Ludyga et al., 2016).

9 It has to be noted that, while some research supports the notion that specific physical or cognitive trainings generalize beyond the specific functions they presumably address (Jaeggi, Buschkuehl, Jonides, & Perrig, 2008; Sherrington, Tiedemann, Fairhall, Close, & Lord, 2011), there is also an accumulating body of research that indicates that such effects are task-specific and do not generalize (Giboin, Gruber, & Kramer, 2015; Giboin et al., 2019, Melby-Lervag & Hulme, 2013).