Risk matrix literature review

The risk matrix, originally seen as a simple risk analysis technique, has recently excited controversy with different authors promoting its use and others suggesting it has little value. The purpose of this section is to:

- explore the advantages and disadvantages of the risk matrix
- show how it can be designed and used to give reasonable results
- suggest training and written instructions needed to promote such results or lead to useful, relevant results.

Origins, definitions and applications of the risk matrix

Ale (2007) credited Napoleon with use of a risk matrix based on the likelihood of consequences. Hussey (1978) described a two-dimensional directional policy matrix to aid decision-making, and Schleier & Peterson (2010) reviewed use of the Entomological Operational Risk Assessment (EORA) by the US military to estimate risks posed by zoonotic pathogens carried by insects.

Ball & Watt (2013) gave a critical review of the risk matrix use in non-occupational and low-risk occupational settings. The risk matrix was applied in the safety sciences in the late 1980s in the UK, with simple versions being described during a UK Institution of Occupational Health and Safety conference in Belfast[1] and by Moore (1997).

Relevant definitions

When developing and using a risk matrix it is important to keep in mind the definition of risk. As used here, risk is defined in ISO31000 *Risk management: principles and guidelines* as the:

> *effect of uncertainty on objectives.*
>
> > *Note 1: An effect is a deviation from the expected (positive and/or negative).*
> >
> > *Note 2: objectives can have different aspects (such as financial, health and safety, and environmental goals and can apply at different levels (such as strategic, organisation-wide, project, product and process).*
> >
> > *Note 3: risk is often characterised by reference to potential events and consequences or a combination of these.*
> >
> > *Note 4: risk is often expressed in terms of a combination of the consequences of an event (including changes in circumstances) and the associated likelihood of occurrence.*
> >
> > *Note 5: uncertainty is the state, even partial, of deficiency of information related to, understanding or knowledge of an event, its consequences or their likelihood.*

From this it can be seen that a risk assessment is a description of some uncertain future state. Uncertainty also arises, in part, because of limitations in risk assessment techniques, including the risk matrix.

Cook (2008) described the risk matrix as a tool or technique for assigning a level of risk to the outcomes of an event. Other authors have described it as a two-dimensional tool with scales for the consequences of an event and their associated likelihood or probability but, more formally, ISO Guide 73 (ISO, 2009) defines risk matrix as a:

> *tool for ranking and displaying risks by defining ranges for consequence and likelihood.*

The output from a risk matrix will be the level of risk, defined in ISO Guide 73 as the:

> *magnitude of a risk or combination of risks, expressed in terms of the combination of consequences and their likelihood.*

[1] Personal communication, Hani Raffart, 1991

Summary of the advantages and disadvantages of the risk matrix

The following table summarises the main advantages and disadvantages of the risk matrix and suggests some solutions for the disadvantages.

It is important to keep in mind that even the best designed risk matrix can have severe limitations and may mislead decision-makers who rely on it without any supporting data from other risk analysis techniques.

Table 1. Advantages and disadvantages of the risk matrix

Main advantages of the risk matrix	Main disadvantages of the risk matrix	Solutions to disadvantages
Provides an easily understood representation of different levels of risk	Yields a subjective, not an objective, analysis of risk	Guidance material to help reduce subjectivity; training in the use of the matrix and peer review of the findings
Can be compiled relatively quickly	May rely on subjective or incomplete analysis of incomplete historical data	
Aids decision-making about or under risk	Lacks granularity (eg, a five-point scale cannot represent a wide range of consequences or likelihoods)	Consider using a log scale if users understand the maths
Promotes discussion in risk workshops.	Often uses the same words for the level of risk as for the consequences	Design of the matrix, including consultation and user group testing
Enables decision-makers to focus on the highest priority risks with some consistency	Often uses uncertain, opaque, ambiguous or obscure descriptions	
Enables quick ranking and comparison of risks for attention	Often designed without reference to the risk profile of the organisation or risks being analysed	
Can be used to help develop a common understanding of levels of risk across an organisation.	The worst case high consequence/low likelihood level of risk with controls in place is greater than tolerable	
Enables the combination of consequences and their likelihood to be represented graphically	Consequence scales may not include known consequences in an organisation	
	No or inadequate alignment between risks with different consequences (eg, financial, safety, reputational)	Design of the matrix, including consultation and user group testing. For safety-related risks, consider using "value of a statistical life" data
	May show a risk as tolerable at the corporate level but intolerable at a departmental or unit level	Design of the matrix, including consultation and user group testing. Guidance to users and decision makers
	Inappropriate numerical scoring of cells	Design of the matrix, including consultation and user group testing. Use qualitative descriptors of the level of risk (eg, extreme, high, low)
	May tempt users to under- or over-state the consequences or their likelihood, resulting in incorrect analysis of the level of risk Some users may not consider all available data	Guidance material to help reduce subjectivity; training in the use of the matrix and peer review of the findings Train the trainers and moderate their work
	May use historical data that includes conditions that were unnoticed, unrecorded or unrepeatable, leading to misinterpretation and misjudgement	Competent review of data and analyses
	Often only shows negative consequences	Design and use of a 10x5 matrix
	Narrow use of risk matrix (eg, only analyse risks with controls)	Analyse risks with no controls, then with current controls, then, if necessary, if treatment was applied

Sources: (Cox, 2008; Franks et al, 2002; Hillson & Simon, 2007; Julian, 2011; Middleton & Franks, 2001; SA, 2012)

Details of advantages, disadvantages and some related issues

Failure to understand scales

Bahill & Smith (2009) discussed use of a frequency/severity graph and showed how it could use linear scales or log scales. They also showed how essential it is to use appropriate risk frequency and severity scales to avoid misrepresenting the level of risk or giving a false picture to decision-makers.

Ball & Watt (2013) considered it difficult to define qualitative scales accurately and unambiguously but problems can arise if numerical matrix scales are not linear (eg, they are logarithmic) and users are unfamiliar with such scales.

Safety-related scales may appear logical and linear but actually conceal non-linear and potentially illogical categories. For example, a minor injury may subsequently deteriorate and then be described as a major injury. When dealing with human health or safety, matrices rarely consider variations in vulnerability of exposed people such as age, gender, pre-existing health factors, height or fatigue (Duijm, 2015).

IEC/ISO31010 paragraph B29.6 (IEC/ISO, 2009) notes that "it is difficult to define scales unambiguously" and "use is very subjective and there tends to be significant variation between raters".

Other problems with scales used in risk matrices include (Cook, 2008):

- the likelihoods under consideration may be so small as to be indistinguishable
- severity and likelihood categories may be very broad.

Errors in design

The non-linearity of points of equal risk may not align themselves with cells or the boundaries between cells in a matrix (Pickering & Cowley, 2010). This problem cannot be overcome by changing the position of gridlines or the number of rows or columns. As a result, the level of risk may be over- or under-estimated.

Figure 1 shows an example of a risk matrix. It is used in this paper to help illustrate design errors and how some errors can be overcome.

In this matrix, as recommended by Cox (2008), small increases in likelihood or consequence will not cause a jump in the level of risk from, for example, low to high without going through the medium category.

Figure 1. Example of a consequence/likelihood matrix

Likelihood ↑		A	B	C	D	E
	V	Medium 5	High 10	High 15	Extreme 20	Extreme 25
	IV	Medium 4	Medium 8	High 12	Extreme 16	Extreme 20
	III	Low 3	Medium 6	Medium 9	High 12	Extreme 15
	II	Low 2	Low 4	Medium 6	High 8	Extreme 10
	I	Negligible 1	Low 2	Medium 3	High 4	Extreme 5
		Consequences →				

Cox argued a matrix should not use too many colours or labels to represent levels of risk. Three colours (eg, red, yellow and green) or levels of risk seemed a minimum and five a maximum (thus, Figure 1 is at the limits of reliable matrix design). This point is further discussed later.

Use of a risk matrix

Experience suggests that many people will use a risk matrix to analyse the level of risk:

- either as if there were no controls
- or as if the controls were perfect.

This overlooks that some controls are likely to be in place but are not likely to be perfect.

A matrix can be used to analyse three levels of a given risk. The analyses should:

- not take into account any controls under the direct influence of the organisation that were absent or had failed; this may show a high consequence with a high likelihood (level of risk 1)
- take into account the effectiveness of the controls "as is"; this often will show a lower consequence or likelihood of that consequence (level of risk 2)
- consider the plausible consequences or their likelihood if practicable ways of modifying the risk were introduced (level of risk 3).

This approach will show the importance of the controls or the possible treatment. Controls shown to change the level of risk from, for example, extreme to low might be regarded as key controls that should be regularly monitored by management. Similarly, risk treatments that, if implemented, will change a safety-related risk from high to negligible should be monitored from approval to completion.

Worst-case high consequence/low likelihood with current controls

Good practice design of a risk matrix should result in the level of risk from the combination of the worst consequence and lowest likelihood of that consequence being tolerable to top management. This must take account of the effectiveness of the current controls or the level of risk after treatment.

For example, using Figure 1 to estimate level of risk 1 (ie, no controls), the worst consequence might be "E". If the likelihood of that consequence was "II" the level of risk would be "extreme". However, after considering the effectiveness of the current controls the worst consequence might be analysed as "D" with a likelihood of "I", showing a level of risk 2 of "high". Some treatment options might be identified that would bring the worst consequence to "C" with a likelihood of "I" with a forecast level of risk 3 of "medium".

Inappropriate quantification

Often, attempts are made to quantify a matrix by allocating scores to the consequence and likelihood scales. This might be done on a linear scale (1, 2, 3, 4, 5) resulting in a range of scores in Figure 1 from 1 (1x1 = 1 in cell IA) to 25 (5x5 = 25 in cell VE). However, these results would not match the descriptions used for the cells: the cell in the bottom right-hand corner scores 1x5 = 5 but is rated as extreme.

In an experiment Ball & Watt (2013) used a symmetrically scored 5x5 matrix and asked 50 students to rate the level of risk in three different, non-occupational settings. The students rated the risks with a wide range of scores between 2 and 25, but with considerable clustering in the range 10-15. The scattering and clustering of results suggested quantified matrices may not give sufficiently reliable results for decision-making.

In their experiment, Ball & Watt found the lack of information about the risks suggested a matrix should only be used after earlier application of risk assessment techniques intended to gather qualitative information about the nature of a risk. They also discussed issues associated with different risk perceptions of stakeholders, each of whom may have different objectives affected by the risk and different experiences of risk.

Ball & Watt thought that risk assessors rarely complete a risk assessment by developing and implementing treatments for a risk evaluated as unacceptable; further, they often do not reassess the risk or receive feedback on their performance and, as a result, their performance may not improve (Tetlock, 2006; Tetlock & Gardner, 2015).

Ball & Watt discussed Kahneman's System 1 and System 2 model for thinking about risk (Kahneman, 2011) and suggested many people will use a matrix with fast, intuitive thinking (ie, System 1). This would lead to Kahneman's "What you see is all there is" conclusion with no in-depth analysis, overconfidence and inaccurate framing of the risk. System 2 thinking (slower, more considered

thinking using best-available evidence) may result in frustration because a risk matrix lacks sufficient detail or is not supported by information derived from a preceding risk analysis.

A matrix designer might attempt to apply asymmetrical level-of-risk values. For this to be a valid approach the designer would need a substantial body of data on which to base the chosen values. Such a database would take time to build and might use, for example:

- historical data related to an environment that since has changed, so giving false results
- incomplete data, giving rise to uncertainty
- data under-reported by those responsible for an adverse loss, giving rise to uncertainty about "washed" data
- data reported by people on the "winning team", giving rise to uncertainty due to over-stated results
- use of data from the context of one risk that is not relevant to the context of another risk.

"Layering"

Further problems arise when designers attempt to reduce the apparent uncertainties in a matrix by "layering" either qualitative or quantitative user questions before use of a matrix. These questions may be subject to framing errors and designer bias, so introducing hidden uncertainties, such as the "probability of a probability".

Subjectivity in the use of matrices

The use of risk matrices involves "subjective and arbitrary judgements", making any risk assessment based on a matrix alone questionable (Bluff & Johnstone, 2004).

The causes of subjectivity are many and diverse but it is well-known that people have a tendency to overestimate small probabilities and underestimate large probabilities, leading to clumping of risk assessments towards the centre of a matrix (as found by Ball & Watt).

All risk assessments rely on the subjectivity and "professional judgement" of the risk assessors, who may lack training in risk assessments. This level of subjectivity may then exacerbate the subjectivity designed into a matrix.

Behavioural effects

Smith, Siefert, & Drain (2009) carried out a cross-disciplinary examination of the risk matrix and showed it is prone to design errors arising from cognitive biases in designers. They used Prospect Theory (Kahneman & Tversky, 1979) to show how framing effects can distort placement of matrix reference points (boundaries between cells).

They also showed matrix-users will tend to place consequence/likelihood combinations on a line drawn diagonally from bottom left to top right in Figure 1. This potentially results in the bottom right cell in Figure 1 (high consequence, low likelihood) being under-used.

There might also be an exaggeration of potential harm or loss if people have a personal interest in the outcome. Conversely, understating such potential may occur if the resultant risk assessment must then be reported up the management chain, resulting in unwelcome management attention.

Evans (2012) argued that individual people have different risk tolerances. This can further distort how a matrix is used: people with low risk tolerance will over-rate a risk while those with higher risk tolerance will under-rate it.

For all these reasons, an agreed organisational risk tolerance must be identified before a risk matrix is used.

Lack of data or wrong data

Wintle, Kenzie, Amphlett, & Smalley (2001) thought that, used as part of a more detailed risk analysis, matrices could help duty holders to identify items of plant presenting the greatest risks arising from failure. They also noted that (p. 49):

> *High consequence/low frequency failures are difficult to quantify since the events may never have occurred and are by their nature extremely rare. Often a cut-off is imposed at a level of consequence above which all levels of risks are considered significant. In these cases, an*

> *appropriate level of inspection during service is prudent to provide continuing assurance that the assumptions of the analysis remain valid.*
>
> *Particular difficulties arise for new equipment or processes where there is a lack of data needed to assess the risk (eg, lack of design or materials data, operational or inspection history, or limited analysis of consequences). More frequent inspections may be necessary at the start of life to demonstrate integrity under operating conditions".*

Records may show a specified type of event has known consequences or frequencies, but if matrix designers are unaware of that data (an "unknown known"(Aven, 2015; Zizek, 2004)), incorrect design of the consequence and likelihood scales will result (Smith et al, 2009). Similarly, matrix users may lack necessary knowledge of events, their consequences and the likelihood of the consequences.

In the absence of historical data to aid use of a matrix it may be necessary to make assumptions, but such assumptions may be biased by the interests or experience of the risk assessors (Ball & Watt, 2013). A further problem may arise when there is inadequate data for a risk matrix-based analysis and an expert is used to provide an opinion: if there is insufficient data, it is open to question what the expert is basing their opinion on (Redmill, 2002).

Bahill & Smith (2009) similarly argued:

> *The data used in a risk analysis have different degrees of uncertainty; some of the values are known with precision, others are wild guesses; however, in addition to uncertainty, all data have opportunity for errors.*

This is a key criticism of risk matrices: they are often perceived as a scientific tool because they contain numbers, even though the input numbers contain unstated uncertainties – even "wild guesses". Some of those uncertainties may be "back of an envelope" calculations, estimates or guesses made when the matrix was being developed. It therefore is crucial the designer of a matrix states the assumptions and uncertainties in a matrix.

It is often claimed that risk matrices are simple and transparent in use. However, the apparent rational – even scientific – appearance of matrices can disguise development problems such as those outlined above and give rise to deception of risk assessors and decision-makers.

Language used

The language used in categories may not be immediately clear to a risk assessor or decision-maker and associated descriptions may require interpretation, interpolation or even guesses as to meaning, leading to incorrect risk analyses (Duijm, 2015; Kent, 2007).

Variability in interpretation of language (eg, consequence or likelihood descriptions) can be considerable (Christensen et al, 2003; Hayes, 2010; Kent, 2007) and is difficult to overcome. While some judgement needs to be left to risk assessors, this issue increases uncertainty about results from use of a risk matrix.

Errors in use

A risk matrix might show that an organisation has many low consequence/high likelihood risks which may then be downplayed by management, leading to an organisational culture of "no real harm done" and, ultimately, erosion of the effectiveness of an organisation.

Authors of the articles reviewed for this paper consistently refer to the use of the risk matrix as a tool for ranking risks for urgency of attention. For example, guidance in HB 89: 2012 *Risk management – Risk assessment techniques* (SA, 2012) describes the risk matrix as a screening tool, while Donoghue (2001) describes the design of qualitative and semi-quantitative matrices to aid operational decision-making after walk-through inspections.

It needs to be made clear to matrix users that their risk matrix is an aid to analysing the likelihood of specified consequences, not the likelihood of the event giving rise to those consequences.

Assessors are untrained or lack competence

Ball & Watt showed that users in non-occupational or low-risk occupational settings might be inexperienced in risk assessments using a risk matrix, giving rise to results that are potentially inconsistent between users or so inaccurate as to give rise to ridicule.

The level of detail ("granularity") used in the consequence and likelihood scales may be inadequate to do more than give an indication of the level of a risk. For example:

- the boundary between two financial consequences might be $100,000; inexperienced risk assessors may be tempted to analyse a negative consequence as less than $100,000 or estimate a positive consequence as greater than $100,000
- when considering the likelihood of such consequences, inexperienced risk assessors may misremember or never have heard of such a negative consequence or be anxious that a project goes ahead.

Such inaccuracies might place a risk in any group of four contiguous cells in Figure 1. Depending on the selected consequence and likelihood points, this could give levels of risk of:

- extreme, high, or medium
- high or medium
- medium or low
- low or negligible.

Accuracy and reliability of the results

Cox (2008), in an exhaustive review of matrices, concluded that his theoretical results generally showed quantitative and semi-quantitative risk matrices have limited ability to correctly reproduce the risk ratings implied by quantitative models – a key theoretical finding.

Risk matrices may be activity- or system-centric and so may not explicitly consider the risks to an individual (Wilkinson & David, 2009). They may also encourage a focus on small components of a system or project ("salami-slicing") and so not show the overall level of risk.

A contract research report for the UK HSE (Wintle et al, 2001) reviewed best practice for risk-based inspection as part of plant integrity management. At page 48 the authors noted:

> *Risk matrices are a useful means of graphically presenting the results from qualitative risk analyses of many items of equipment. Risk matrices should, however, not be taken too literally since the scale of the axes is only indicative.*

Possible change in guidance

IEC/ISO31010: 2009 *Risk assessment techniques* is currently undergoing revision. One of the lead authors has said the committee believes "the matrix is not an analysis technique but a display technique because you must analyse consequences and likelihood first then just display it on the matrix"[2]. This is a subtle but important distinction.

Summary

Risk matrices are an overrated way of analysing the level of risk. While risk matrices apparently provide a simple mechanism for analysing the level of individual risks they too often are poorly designed or incorrectly used.

Any risk matrix represents an attempt to fit a wide range of dynamic risk levels into a narrow range of predefined cells. This requires qualitative judgement in what may be assumed to be a quantitative process.

A risk matrix should be used as one of a set of risk analysis techniques. Used alone, risk matrices are likely to give misleading, even wrong, results leading to false certainties or inappropriate allocation of resources.

The results from use of a matrix should be treated as an indicative comparison between risks, to help distinguish those requiring urgent attention.

If a risk matrix is to be used, it should be

- simple, and designed using data relevant to the organisation
- used with a clear understanding of the nature of a risk.

The limitations of any risk matrix must be understood by risk assessors and decision-makers.

[2] Personal communication from Jean Cross, 1 December 2015.

References

Ale, B. (2007). *Risk is of all time*. Working Paper Delft.

Aven, T. (2015). Implications of black swans to the foundations and practice of risk assessment and management. *Reliability Engineering & System Safety, 134*(0), 83-91.

Bahill, A. T., & Smith, E. D. (2009). An Industry Standard Risk Analysis Technique. *Engineering Management Journal, 21*(4), 16-29.

Ball, D. J., & Watt, J. (2013). Further Thoughts on the Utility of Risk Matrices. *Risk Analysis*, n/a-n/a.

Bluff, L., & Johnstone, R. (2004). *The Relationship Between 'Reasonably Practicable' and Risk Management Regulation*. Working Paper Canberra: Australian National University. Retrieved from http://ohs.anu.edu.au/, 15 July 2010.

Christensen, F. M., Andersen, O., Duijm, N. J., & Harremoës, P. (2003). Risk terminology—a platform for common understanding and better communication. *Journal of Hazardous Materials*, 181-203.

Cook, R. (2008). *Simplifying the creation and use of the risk matrix*. Paper presented at the Improvements in System Safety conference, Bristol organised by

Cox, L. (2008). What's Wrong with Risk Matrices? *Risk Analysis, 28*(2), 497-512.

Donoghue, A. M. (2001). The design of hazard risk assessment matrices for ranking occupational health risks and their application in mining and minerals processing. *Occupational Medicine-Oxford, 51*(2), 118-123.

Duijm, N. J. (2015). Recommendations on the use and design of risk matrices. *Safety Science, 76*, 21-31.

Evans, D. (2012). *Risk intelligence: how to live with uncertainty*. London: Atlantic Books.

Franks, A., Whitehead, R., Crossthwaite, P., & Smail, L. (2002). *Application of QRA in operational safety issues*. Research Report RR 025 for Health and Safety Executive; Sudbury: HSE Books. Retrieved from http://www.hse.gov.uk/research/, 12 January 2003.

Hayes, J. (2010). *Safety Decision Making – Drawing a Line in the Sand*. Working Paper 74 Canberra: Australian National University. Retrieved from http://www.http://ohs.anu.edu.au/, 15 July 2010.

Hillson, D., & Simon, P. (2007). *Practical Project Risk Management: the ATOM methodology*. Vienna, VA: Management Concepts.

Hussey, D. E. (1978). Portfolio Analysis: Practical Experience with the Directional Policy Matrix. *Long Range Planning, 11*(4), 2-8.

IEC/ISO31010:2009. Risk Management - Risk Assessment Techniques. Geneva: International Standards Organisation

ISOGuide 73:2009. Risk management - vocabulary - guidelines for use in standards. Geneva: International Standards Organisation

Julian, T. (2011). What's right with risk matrices? An great tool for managers. Retrieved 26 July 2012, from http://31000risk.wordpress.com/article/what-s-right-with-risk-matrices-3dksezemjiq54-4/

Kahneman, D. (2011). *Thinking fast and slow*. London: Penguin Books Ltd.

Kahneman, D., & Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica, 47*(2), 263-291.

Kent, S. (2007). Words of Estimative Probability. Retrieved 28 December, 2012, from https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/sherman-kent-and-the-board-of-national-estimates-collected-essays/6words.html

Middleton, M., & Franks, A. (2001). Using Risk Matrices. *Chemical Engineer*(723).

Moore, D. (1997). The Use of a Ranking Matrix and Recommendation Prioritization System for Process Hazard Analysis Studies. *Process Safety Progress, 16*(2), 83-85.

Pickering, A., & Cowley, S. (2010). Risk Matrices: implied accuracy and false assumptions. *Journal of Health & Safety Research & Practice, 2*(1), 9-16.

Redmill, F. (2002). Exploring subjectivity in hazard analysis. Retrieved from http://homepages.cs.ncl.ac.uk/felix.redmill/publications/5E.Exploring Subjectivity.pdf, 5 October 2015.

SAHB 89:2012. Risk management – Guidelines on risk assessment techniques. Sydney: Standards Australia

Schleier, J. J., & Peterson, R. K. D. (2010). Limitations of the Entomological Operational Risk Assessment Using Probabilistic and Deterministic Analyses. *Military Medicine, 175*(8), 594-598.

Smith, E. D., Siefert, W. T., & Drain, D. (2009). Risk matrix input data biases. *Systems Engineering, 12*(4), 344-360.

Tetlock, P. (2006). *Expert political judgment : how good is it? how can we know?* Princeton, NJ: Princeton University Press.

Tetlock, P., & Gardner, D. (2015). *Superforecasting: The art and science of forecasting*: Random House.

Wilkinson, G., & David, R. (2009). Back to Basics: Risk Matrices and ALARP. In Dale, C. & Anderson, T. (Eds.), *Safety-Critical Systems: Problems, Process and Practice* (chapter 11, pp. 179-182): Springer London.

Wintle, J., Kenzie, B., Amphlett, G., & Smalley, S. (2001). *Best practice for risk based inspection as a part of plant integrity management*. Research Report CRR 363 for Health and Safety Executive; Sudbury: HSE Books. Retrieved from http://www.hse.gov.uk/research/, 12 January 2004.

Zizek, S. (2004). What Rumsfeld Doesnt Know That He Knows About Abu Ghraib. *In These Times*. Retrieved from http://www.inthesetimes.com/article/747/, 24 January 2013.